

BOW TIES IN THE SKY II: SEARCHING FOR GAMMA-RAY HALOS IN THE *Fermi* SKY USING ANISOTROPY

PAUL TIEDE¹, AVERY E. BRODERICK^{1,2}, MOHAMAD SHALABY^{1,2,3}, CHRISTOPH PFROMMER⁴,
EWALD PUCHWEIN⁵, PHILIP CHANG⁶, AND ASTRID LAMBERTS⁷

¹ Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON, N2L 3G1, Canada

² Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON, N2L 2Y5, Canada

³ Department of Physics, Faculty of Science, Cairo University, Giza 12613, Egypt

⁴ Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

⁵ Institute of Astronomy and Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

⁶ Department of Physics, University of Wisconsin-Milwaukee, 3135 North Maryland Avenue, Milwaukee, WI 53211, USA

⁷ TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

Draft version February 10, 2017

ABSTRACT

Many-degree-scale gamma-ray halos are expected to surround extragalactic high-energy gamma ray sources. These arise from the inverse Compton emission of an intergalactic population of relativistic electron/positron pairs generated by the annihilation of $\gtrsim 100$ GeV gamma rays on the extragalactic background light. These are typically anisotropic due to the jetted structure from which they originate or the presence of intergalactic magnetic fields. Here we propose a novel method for detecting these inverse-Compton gamma-ray halos based upon this anisotropic structure. Specifically, we show that by stacking suitably defined angular power spectra instead of images it is possible to robustly detect gamma-ray halos with existing *Fermi* Large Area Telescope (LAT) observations for a broad class of intergalactic magnetic fields. Importantly, these are largely insensitive to systematic uncertainties within the LAT instrumental response or associated with contaminating astronomical sources.

Keywords: BL Lacertae objects: general – gamma rays: general – radiation mechanisms: non-thermal – gamma rays: diffuse background – infrared: diffuse background – plasmas

1. INTRODUCTION

The extragalactic gamma-ray sky at TeV energies is dominated by blazars, a subclass of active galactic nuclei (AGNs) with powerful relativistic outflows directed at us. The relativistic jets are powered by accretion onto a central nucleus, presumably a supermassive black hole. In the unified picture of AGNs, their emission properties depend on the orientation of the AGN relative to the line of sight (Urry & Padovani 1995). There exist two categories of AGNs that differ in their accretion mode and in the physical processes that dominate their emission.

- *Thermal/disk-dominated AGNs.* Infalling matter assembles in a thin disk and radiates thermal emission with a range of temperatures. This emission is then Comptonized by a hot corona above the disk to produce power-law X-ray emission, defining the class of quasars or Seyfert galaxies.
- *Non-thermal/jet-dominated AGNs.* Highly energetic electrons that have been accelerated in the relativistic jet interact with the jet magnetic field and emit synchrotron radiation from the radio to X-ray regime. In addition, the same population of electrons can Compton up-scatter seed photons that are either provided by the synchrotron radiation itself or by an external photon field into the gamma-ray regime. Hence, the broadband spectral energy distribution of these objects is characterized by two peaks. This defines the class of radio-loud AGNs which can furthermore be subdivided into blazars (with the line of sight intersecting the jet opening angle) and non-aligned non-thermal dominated AGNs.

As a result, the gamma-ray emission of blazars benefits from the relativistic Doppler boosting, shifting the upper end of the gamma-ray emission into the GeV/TeV energy regime. Blazars exhibit a continuous sequence: their luminosity anti-correlates with the peak energy of their synchrotron spectrum, i.e., the objects emitting very high-energy gamma rays (VHEGRs) at TeV energies have the lowest intrinsic luminosity (e.g., Fossati et al. 1998; Ghisellini et al. 1998).

The main observational representatives of both AGN classes, quasars and radio galaxies, exhibit a strong redshift evolution with a steeply rising comoving luminosity density up to a redshift $z \approx 2$ and a decline thereafter (Hopkins et al. 2007). In contrast, we can only observe nearby TeV blazars that typically reach out to redshifts of $z \lesssim 0.3$ (Wakely & Horan 2008[†]). The reason for this apparent contradiction lies in the low luminosity of TeV blazars and the finite mean free path of TeV photons as they propagate through space (Ackermann et al. 2012b; Domínguez et al. 2013), precluding the detection of high-redshift blazars (if they exist). The opacity of the Universe to TeV photons is due to the annihilation and pair production of TeV photons of energy E_γ on the extragalactic background light (Gould & Schröder 1967; Salamon & Stecker 1998). The mean free path of these VHEGRs is

$$D_{\text{pp}}(E_\gamma, z) \approx 5.8 \left(\frac{1+z}{2} \right)^{-\zeta} \left(\frac{E_\gamma}{6 \text{ TeV}} \right)^{-1} \text{ Mpc}, \quad (1)$$

where $\zeta = 4.5$ for $z < 1$ and $\zeta = 0$ for $z \geq 1$ (Kneiske et al. 2004; Neronov & Semikoz 2009). Momentum conservation ensures that the pairs propagate essentially in the same direction of the parent TeV photon and energy conservation implies a pair energy of $E_{e^\pm} \approx E_\gamma/2$.

[†] See <http://tevcat.uchicago.edu>, Catalog Version 3.4

The resulting ultra-relativistic pairs of electrons and positrons are commonly assumed to lose energy primarily through inverse Compton (IC) scattering with photons of the cosmic microwave background (CMB), cascading the original TeV emission down to (multi-)GeV energies on a mean free path for the scattering process of

$$D_{\text{IC}} = \frac{3m_e^2 c^4}{4\sigma_T u_{\text{CMB}} E_{e^\pm}} \approx 0.1 \left(\frac{E_{e^\pm}}{3 \text{ TeV}} \right)^{-1} (1+z)^{-4} \text{ Mpc}, \quad (2)$$

where $m_e c^2$ is the electron rest mass energy, σ_T is the Thompson cross section, and u_{CMB} is the CMB energy density.

However, the inverse Compton cascaded (ICC) multi-GeV emission has not been observed by the *Fermi* Large Area Telescope (LAT), indicating that some additional physics needs to be considered (see, e.g., Neronov & Vovk 2010; Tavecchio et al. 2010, 2011; Dermer et al. 2011; Taylor et al. 2011; Takahashi et al. 2012; Dolag et al. 2011; H. E. S. S. Collaboration et al. 2014; Prokhorov & Moraghan 2016). The presence of intergalactic magnetic fields (IGMFs) would deflect the beam of e^+/e^- pairs with a Larmor radius of

$$r_L = \frac{E_{e^\pm}}{eB} \approx 30 \left(\frac{E_{e^\pm}}{3 \text{ TeV}} \right) \left(\frac{B}{10^{-16} \text{ G}} \right)^{-1} \text{ Mpc} \quad (3)$$

out of our line of sight. This reduces the ICC flux and thus provides a lower limit on the strength B of the IGMF. For the associated ICC photons with energies of

$$E_{\text{IC}} \approx 2\gamma^2 E_{\text{CMB}} \approx 70 \left(\frac{E_{e^\pm}}{3 \text{ TeV}} \right)^2 \frac{E_{\text{CMB}}}{1 \text{ meV}} \text{ GeV} \quad (4)$$

the *Fermi* angular resolution is $\theta \approx 0.2$ or 3×10^{-3} rad (using the $1-\sigma$ containment angle of combined events, see Fig. 57 of Ackermann et al. 2012a). Hence, a deflection of the pairs by an angle $D_{\text{IC}}/r_L > \theta$ implies a lower limit on the IGMF of $B \gtrsim 10^{-16} \text{ G}$ (see e.g., Neronov & Vovk 2010) with important implications for primordial magnetogenesis (for reviews, see Kandus et al. 2011; Durrer & Neronov 2013).

Alternatively, the process of ultra-relativistic pairs propagating through the intergalactic medium can be viewed as two counter-propagating beams that are subject to plasma instabilities, and, in particular, to the oblique instability (Broderick et al. 2012). The linear growth rate of the oblique instability is larger than the cooling rate due to IC scattering of the pairs (Broderick et al. 2012; Schlickeiser et al. 2012). If this dominance of the instability growth rate carries over to the regime of non-linear saturation (Chang et al. 2014; Schlickeiser et al. 2013, but see also Sironi & Giannios 2014; Miniati & Elyiv 2012), this causes the kinetic energy of the pairs to be transferred to the unstable electromagnetic modes in the background plasma. This energy should eventually be dissipated, heating the intergalactic medium (Chang et al. 2012) with potentially far-reaching implications for the Lyman- α forest (Puchwein et al. 2012; Lamberts et al. 2015), non-linear structure formation (Pfrommer et al. 2012), and for the blazar luminosity function and the extragalactic gamma-ray background (Broderick et al. 2014a,b). Estimates suggest that ~ 300 yr after the onset of TeV emission, the pair beam density has grown sufficiently for plasma beam instabilities to dominate its evolution. This would randomize the beam, and potentially suppress the ICC emission on which the IGMF limits are based (rendering these limits dubious, Broderick et al. 2012).

How can we determine the ultimate fate of these pairs? Clearly an unambiguous detection of the deflected pair halo emission would immediately prove the deflection hypothesis. So far, all work has concentrated on measuring excess halo power at large angular scales through stacking analyses of blazar images. However, those resulted in null results (e.g., Ackermann et al. 2013a) despite some earlier claims (Ando & Kusenko 2010) that were subsequently disproven (Neronov et al. 2011) because of the uncertainty of the exact shape and side lobes of *Fermi*'s point-spread function (Ackermann et al. 2012a). A more recent attempt (Chen et al. 2015) utilizes the most recent PSF; nevertheless, it exhibits similar sensitivities to the uncertain instrument response.

Here we report a novel method for extracting signatures of the ICC component that exploits the large degree of anticipated anisotropy (see, e.g., Neronov et al. 2010; Long & Vachaspati 2015; Broderick et al. 2016). This is caused either by the structure of the initial VHEGR jet (if its opening angle is not intersecting the line of sight) or by the fact that both electrons and positrons are produced by the VHEGR annihilation on the EBL and deflected in opposite directions by an IGMF. The details of the structure depends on both the mechanism and properties of the IGM. However, in both cases they can produce halos that are dominated by bi-lobed features. Such features are generic for the gamma-ray bright blazars observed by *Fermi* (Broderick et al. 2016).

In principle, this angular asymmetry can improve our ability to detect ICC halos in two different ways. First, the anisotropy implies a larger surface brightness for the ICC emission, and therefore *if we can properly orient the images* any excesses can be detected with a higher significance over the putative symmetric backgrounds. In practice, we do this not by rotating and stacking images, which proves to be unfeasible due to the inability of determining the orientation of the deflecting magnetic field, but rather by constructing and stacking orientation-independent measures of the angular anisotropy in the gamma-ray sky about individual sources. In particular, we propose to exploit the pair halo anisotropy by computing suitably defined angular power spectra and stacking those. This approach is routinely applied in cosmological data analyses to determine cosmological parameters (e.g., from the CMB anisotropies) – the key idea consists of averaging over the (unknown) orientation of the phases and accumulating signal in the multipoles of the dominant angular structures.²

Second, averaged over the long live-time of the *Fermi* LAT the resulting PSF is very nearly isotropic. Small residual angular structure in the PSF arising from the cubical geometry of the LAT enters at the hexadecapole ($m = 4$) order. As a result, unlike searches that focus on radial excesses, schemes that exploit the nearly quadrupolar nature of the halos are easily distinguished from systematic effects due to the PSF. While this paper presents the detailed methodology and addresses the (known) systematics, we apply our blind experiment to *Fermi* LAT data in a companion letter (Tiede et al. 2016).

Of particular practical importance is the fact that the *Fermi* data set is not easily repeatable. As a result we have expended considerable effort to predict the anticipated anisotropy signal from various ICC halo models, design and optimize the con-

² A similar strategy is described in Duplessis & Vachaspati (2017), where the utility of the Q statistic is explored, a quantity that is closely related to what we describe as the quadrupolar power. A key distinction between the approaches presented here and in Duplessis & Vachaspati (2017) is that we also construct and utilize the power at the other multipoles, providing an independent characterization of potential systematic errors.

struction of ICC halo diagnostics prior to analyzing the *Fermi* data. Here we report the results of this optimization investigation, i.e., the manner in which mock realizations were produced (Section 3), a cursory description of the characteristics of the *Fermi* hard gamma-ray blazar sample (Section 4), investigation of potential confounding features (Section 5), and the confidence levels with which the various ICC models could be excluded given a null result (Section 6). It is in this last step that the dividends of having done so become manifest – by training the analysis on simulated data (only weakly informed by the gross properties of the *Fermi* sample) we ensure that the resulting conclusions are governed by *a priori* statistics, and are therefore well understood.

2. METHOD OVERVIEW

Before describing the creation of physical realistic ICC halos, statistical measures of the anisotropy, and estimates of its detectability, we will begin with a summary of the key ideas underlying the signal we hope to find. This is predicated on the standard picture of the ICC halo formation described in Section 1: VHEGRs emitted from an AGN travel cosmological distances prior to generating energetic electron-positron pairs on the EBL which then inverse Compton up-scatter the CMB to GeV energies. However, for two independent reasons, these ICC halos are anisotropic.

First, the VHEGRs are originally beamed along the jet axis. This is evidenced by the overwhelming dominance of blazars in the extragalactic gamma-ray AGN sample (Ackermann et al. 2011, 2015b). Because the VHEGR mean free path is long in comparison to the distance traveled during the inverse-Compton cooling time of the resulting pairs this implies that the emission is essentially local, and therefore arises from a pair of conical regions indicated by the radio jet of the source AGN. If the inverse-Compton gamma rays are isotropically emitted, arising, e.g., from a highly tangled IGMF, the spatial structure in the gamma rays generates a resultant structure in the GeV image. This is shown explicitly in the left-hand panel of Figure 1, along with the associated gamma-ray image.

Alternatively, the process of gyration in the IGMF also can impart structure on the image. In the presence of an IGMF that is homogeneous on scales comparable to D_{pp} electrons and positrons will gyrate on fixed trajectories that emit towards an observer only for a subset of initial injection positions. This is still superimposed on the jet structure, resulting in a potentially asymmetric image structure, shown in the right-hand panel of Figure 1. Gamma rays on opposite sides of the original AGN are produced predominantly by different lepton species, i.e., positrons on one side and electrons on the other³.

A toy example that provides many of the key features of the ICC halos we will describe in detail in Section 3, is shown in Figure 2. This includes equal contributions from a uniform background and from a central source, totaling 4000 photons and comparable to a typical bright *Fermi* AGN. In addition, there is an anisotropic halo component (red) containing 10% of the photons in the source drawn from the ad hoc flux distribution indicated by the contours. All components have been convolved with a Gaussian PSF with standard deviation $0''.2$, comparable to the scale of the Pass 8R2_V6 PSF for front-

converted events in the LAT.

In the absence of the component color coding the sub-dominance of the ICC excess makes it difficult to identify directly from a single source. Typically, this is dealt with via stacking multiple images, increasing the statistical significance with which the halo component can be isolated. We therefore show a stacked image of 18 realizations of the same cartoon halo in Figure 2. Because the orientation of the putative ICC halo feature is also randomly varied (corresponding to different source and IGM orientations), in contrast to the single image, the stacked image exhibits nearly no angular structure. Nevertheless, there is a small gamma-ray excess at large angular scales. In Figure 3 this is shown explicitly in comparison to the case when the halo is absent, beginning near angular scales of 1° (where the central source ceases to dominate over the background). The interpretation of the excess is complicated, however, by uncertainty in structure of the large-scale tails of the PSF or the background flux: even marginal modifications of either can absorb the halo signal in its entirety.

Instead we focus on the anisotropic structure of the ICC halo, which presents a unique signature that is difficult to confuse with instrumental response. Explicitly, we construct an angular power spectrum of the surrounding photon positions about the source, defined by:

$$\mathcal{P}_m \equiv \frac{1}{N^2} \left| \sum_j e^{im\theta_j} \right|^2 \quad (5)$$

where θ_j is the polar angle of the j th gamma ray about the image center relative to a fiducial direction and N is the total number of gamma rays. To remove the bulk of the source contribution and eliminate the unresolved structure near the origin we mask the inner regions prior to constructing \mathcal{P}_m ; for illustrative purposes here we simply exclude the inner $0''.6$, though in practice we implement an energy-dependent mask (see Section 5.3). The ICC halos generate a characteristic power spectrum due to their bimodal structure, shown by the red points in Figure 4, that is dominated by $m = 2$ and the even multipoles that follow and qualitatively distinct from most potential image contaminants. In contrast, Poisson noise from a cylindrically symmetric source (e.g., the PSF-convolved central AGN or background) is flat, shown by the black line in Figure 4, and therefore easily distinguished from the anisotropic ICC halos.

More importantly, the \mathcal{P}_m are independent of the orientation of the halo structure. Thus it is possible to stack the \mathcal{P}_m from many sources directly, improving their estimation and thereby improving the significance with which halos may be detected. For example, the \mathcal{P}_m arising from stacking the angular power spectra of the same 18 realizations used to generate Figure 2 is shown explicitly by the blue line in Figure 4. In this the halo structure signal at $m < 10$ is clearly evident in comparison to the Poisson fluctuations that dominate at $m > 15$.

Executing this in practice requires physically realistic halo flux distributions that connect the energy-dependent flux distributions to the underlying physical properties of the VHEGR emission and IGMF and the construction of mock *Fermi* images, to which we now turn our attention.

³ Strictly speaking, the identification of lobe sides with particular leptons does assume that the gyration timescale is long in comparison to the inverse-Compton cooling timescale. Should this be the case today, it will continue to be true at higher redshift. Note that even if the inverse-Compton cooling time is longer than the gyration period, the halo will still exhibit the bimodal structure.

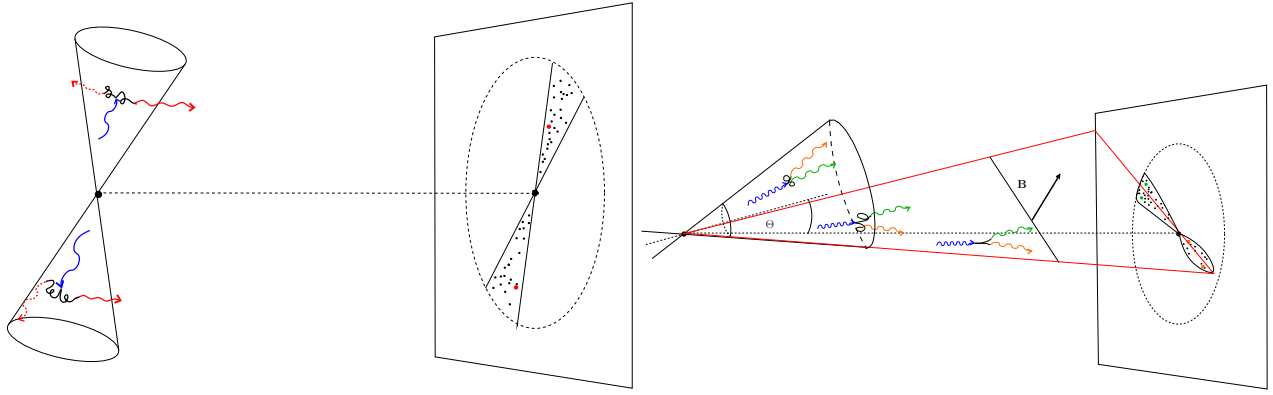


Figure 1. Cartoons of the mechanisms by which anisotropy in the ICC halos is generated, distinguished by the structure of the underlying IGMF. Left: For an IGMF tangled on small scales ($\lambda_B \ll 3$ Mpc) the anisotropy is due to the structure of the gamma-ray jet. Right: For an IGMF that is uniform across the gamma-ray jet (IGMF coherence length $\lambda_B \gg 100$ Mpc) the anisotropy is due to the geometry of the gyrating, relativistic pairs.




Figure 2. Left: Realization of 4000 photons from a toy model that exhibits the anisotropy arising from the mechanisms described in Figure 1. The contribution from the central source, diffuse background, and ICC halo component are shown in purple, green, and red, respectively. The excluded central region is shown in black. Note that despite the clear structure within the ICC halo component it is strongly subdominant to the otherwise isotropic source and background components (contributing roughly 5% of the total photons). Right: Stacked toy model images of 18 sources with arbitrary ICC halo orientations. At this point the structure is effectively erased.

3. GENERATING MOCK IMAGES OF GAMMA-RAY SOURCES

Key to assessing any scheme to detect the ICC halos is the creation of credible theoretical realizations of gamma-ray images of potential sources. While the Section 2 introduced the qualitative reasons to expect anisotropic structure in the ICC halo structure, how to do this quantitatively was presented in Broderick et al. (2016), which we summarize here.

The energies of the gamma rays that comprise the putative ICC halos lie typically below 100 GeV; much higher energy gamma rays are absorbed on the EBL. Below 1 GeV the *Fermi* LAT PSF typically broadens substantially, limiting efforts to find asymmetric features and extending the contaminating influence of bright sources. Between 1 GeV and 100 GeV the *Fermi* LAT response functions only modestly depends on energy. Therefore, we restrict our attention to this energy range.

At the most granular level, *Fermi* images consist of collections of individual photons, numbered in the thousands for a single bright source, each with a reported sky location and energy. Thus, in principle this procedure consists of first, identifying the joint probability distribution of photons from various emission components with a given energy and location, dF/dEd^2x , and second, efficiently drawing a random realization from this, $\{E_j, \mathbf{x}_j\}$. In practice, this is further modified by the *Fermi* LAT response, which primarily impacts the images via the PSF. We consider a three component model comprised of a uniform background, an intrinsic point source, and a putative ICC halo. The former two are well defined and have parameters fixed by *Fermi* directly. Less clear are the ICC halos. Their brightness and morphology depend on the poorly constrained VHEGRs, and thus require some spectral and col-

Table 1
Optimized Source List for a Large-scale, Uniform IGMF

3FGL Name	Common Names	E_p^a (GeV)	Γ_l^b	Γ_h^c	z	R_{img}^d (deg)	$N_{\text{ph},f}^e$ (ph)	B_f^f (ph/deg ²)	$N_{\text{ph},b}^g$ (ph)	B_b^h (ph/deg ²)	B_0^i (G)
3FGL J1104.4+3812	Mkn 421	95.38	1.77	$2.14^{+0.11}_{-0.10}$	0.03	2.0	4999	43.42	4746	75.37	$> 10^{-17}$
3FGL J2347.0+5142	–	1.81	1.69	$1.90^{+0.05}_{-0.04}$	0.044	2.5	3176	144.57	2900	133.47	–
3FGL J1653.9+3945	Mkn 501	236.30	1.72	$2.72^{+0.51}_{-0.37}$	0.034	1.8	2028	56.06	1990	74.14	–
3FGL J2000.0+6509	–	658.00	1.87	$4.25^{+1.38}_{-1.60}$	0.047	3.5	6610	142.40	6257	135.76	–
3FGL J1015.0+4925	–	2.71	1.75	$1.93^{+0.03}_{-0.03}$	0.212	2.5	1797	34.97	1756	38.88	$> 10^{-16}$
3FGL J1444.0-3907	–	11.9	1.67	$2.14^{+0.13}_{-0.09}$	0.065	2.5	2908	124.25	2781	118.81	–
3FGL J0650.7+2503	–	281.60	1.67	$2.29^{+1.64}_{-1.35}$	0.203	2.5	2142	92.84	2069	90.54	–
3FGL J1120.8+4212	–	28.65	1.56	$2.15^{+0.29}_{-0.19}$	0.124	3.0	1123	34.57	1166	35.49	–
3FGL J1442.8+1200	–	0.49	2.69	$1.81^{+0.08}_{-0.07}$	0.163	2.5	921	41.95	998	47.33	–
3FGL J0508.0+6736	–	2.32	1.81	$1.43^{+0.03}_{-0.03}$	0.34	2.0	1927	129.53	1817	129.53	–
3FGL J0303.4-2407	–	1.02	1.78	$1.97^{+0.03}_{-0.03}$	0.26	2.0	984	30.18	1000	41.48	$> 10^{-15}$
3FGL J0543.9-5531	–	0.84	0.46	$1.77^{+0.05}_{-0.04}$	0.273	2.0	1051	73.31	1039	70.55	–
3FGL J1436.8+5639	–	1.05	2.54	$1.79^{+0.10}_{-0.09}$	0.15	2.0	599	39.55	559	40.65	–
3FGL J2329.2+3754	–	2.19	2.19	$1.81^{+0.11}_{-0.08}$	0.264	2.0	1194	83.78	1217	92.32	–
3FGL J0958.6+6534	–	46.64	2.35	$1.35^{+0.68}_{-0.34}$	0.367	3.0	2317	62.36	2118	60.70	–
3FGL J0449.4-4350	–	23.03	1.81	$2.30^{+0.11}_{-0.10}$	0.205	2.0	1653	38.17	1532	41.48	$> 10^{-14}$
3FGL J0757.0+0956	–	59.98	2.25	$1.62^{+0.93}_{-0.61}$	0.27	2.0	775	46.71	740	47.40	–
3FGL J0622.4-2606	–	7.75	2.12	$1.63^{+0.14}_{-0.11}$	0.414	2.0	976	57.74	1001	62.28	–

^a Energy of power spectrum break in GeV.

^b Low-Energy photon spectral index.

^c High-Energy photon spectral index with 1σ error.

^d Radius of image selected in degrees.

^e Number of front converted photons within radius selected.

^f Estimate of number of front converted photons per degree squared from background sources.

^g Number of back converted photons within radius selected.

^h Estimate of number of back converted photons per degree squared from background sources.

ⁱ Values of present-day magnetic field for which the source is in the optimized source list for a large-scale, uniform IGMF (see Section 5.6).

limation model that extends the known *Fermi* properties to TeV energies. This introduces a variety of additional poorly known parameters.

In addition to the intrinsic parameters of the source, the structure of the ICC halo depends critically upon the assumed geometry of the IGMF. Here we consider the two limits described in Section 2: a small-scale, tangled field and a large-scale, uniform field. In principle, these correspond to different assumptions about the IGMF power spectrum. In practice, they imply distinct evolution models for the ultra-relativistic electron/positron pairs following their generation by VHEGR photons from the gamma-ray blazars; the pairs' momenta are rapidly isotropized in the former case, while pairs gyrate around the B-field in the latter case, emitting (toward the observer) only when their momentum is directed toward the observer. We imagine that the general situation lies between these two limits, though for a large range of potential IGMF correlation lengths, λ_B , either will be applicable. The regions where each limit applies in relation to the current constraints on the strength and correlation length of the IGMF are shown in Figure 5.

3.1. Small-scale, Tangled IGMF

The first limit is characterized by a rapid isotropization of the pair momenta. A necessary, though not sufficient, condition for this limit is a strong IGMF, i.e., that the pairs gyrate through $\gtrsim 2\pi$ radians. Hence, the local IGMF strength

$$B \gtrsim 2 \times 10^{-14} \left(\frac{E_{\text{IC}}}{10 \text{ GeV}} \right) (1+z)^4 \text{ G}. \quad (6)$$

Additionally, the IGMF must be dominated by small-scale structures, varying over length scales that permit gyration around a number of axes. Ostensibly, this implies that λ_B must be small in comparison to the typical inverse-Compton cooling length, ≈ 700 kpc (see Equation 2). However, in practice it is sufficient to have isotropization in the statistical sense, i.e., multiple independent domains of locally ordered IGMF within the VHEGR jet. This places a weaker constraint, requiring only that λ_B is small in comparison to the width of the jet, typically of order a few Mpc. These produce ICC halos that for the *Fermi* blazars are characterized by only weak anisotropy. The reason for this is the large foreshortening associated with the gamma-ray blazars that suppresses the angular structure that is dramatic at oblique angles. Nevertheless we consider this case for completeness.

As described in Broderick et al. (2016) the flux of halo photons, shown for a typical realization in Figure 6, is spatially and energy dependent. Here we will ultimately be interested in constructing mock realizations of the *Fermi* sky, and will therefore generate realizations that include an ICC halo component, a direct emission component, and a diffuse background component. As in Broderick et al. (2016) we will assume that the source is a point source, and therefore broadened only via the *Fermi* Pass 8R2_V6 PSF. Furthermore, we will usually assume that the background is locally homogeneous, though will explore departures from this in Section 5. The detailed gamma-ray distribution then depends on the total source luminosity, distance, spectral shape, orientation, jet geometry, and background characteristics. We parameterize these in terms of eight




Figure 3. Radial distribution of gamma-ray surface brightness from the stacked toy model images in Figure 2 with (purple) and without (green) the ICC halo included. This is compared with a model comprised of a central point source and uniform background, convolved with a toy Gaussian PSF with standard deviation 0.2° . While the normalizations for the two components differ slightly from their input values, they are capable of providing a good fit even in the presence of the ICC halos. In any case any disparity is insufficient to convincingly rule out modifications of the large-angle wings of the PSF.




Figure 4. Azimuthal power spectra of a single realization (red squares) and stacked azimuthal power spectra of the 18 realizations (blue circles) of the toy ICC halo model. In these the anisotropic halo component generates a clear signal in even multipoles, beginning with the quadrupole, due to the underlying symmetry of the ICC halo images. For reference, the Poisson noise limit is shown by the black line, to which all models asymptote at large m .




Figure 5. Limits on the strength and correlation length of the IGMF, adapted from Neronov & Vovk (2010). The regions described by the small-scale, tangled IGMF and large-scale, uniform IGMF halo models are highlighted in blue and red, respectively. The lower limit on the magnetic field strength in the small-scale, tangled IGMF halo models arises from the requirement that the pairs execute roughly a full gyration within a cooling time. For comparison, limits arising from the SEDs of 1ES 0229+200 and 1ES 0347-121 are shown by the single- and cross-hatched regions, respectively. Limits arising from direct atomic measurements, Faraday rotation, Big Bang nucleosynthesis, and the cosmic microwave background are also shown (For details on the origins of the limits see Neronov & Semikoz 2009).

quantities:

- The source redshift, z .
- The 1 GeV–100 GeV fluence, N_{35} .
- The low-energy photon spectral index, Γ_l .
- The high-energy photon spectral index, Γ_h .
- The energy of the spectral break, E_p .
- The gamma-ray jet opening angle, θ_{jet} .
- The gamma-ray jet viewing angle, Θ .
- The local background photon density, N_B .

The values of z , N_{35} , Γ_l , and N_B are reported or may be estimated directly for the appropriate sub-sample of *Fermi* AGN. The values of Γ_h , θ_{jet} , and Θ typically are known only for a larger population, and must be constructed from the appropriate distribution for the sources of interest. We defer a discussion of what these are and how the sampling is done until Section 4.

3.2. Large-scale, Uniform IGMF

The second limit is characterized by a uniform IGMF across the extent of the gamma-ray jet. In principle, this requires uniformity on scales of D_{pp} . In practice this is reduced for nearby objects (e.g., Mkn 421) and at high energies; in combination these typically require $\lambda_B \gg 100$ Mpc. Unlike the small-scale, tangled IGMF, there is no condition on the magnetic field strength a priori. However, weak fields necessarily




Figure 6. Components of a typical realization of a gamma-ray map of a *Fermi* AGN with an ICC halo associated with a small-scale, tangled IGMF (top) and a large-scale, uniform IGMF with strength 10^{-15} G (bottom). Each has the background (left), source (center-left), halo (center-right), and combined (right) photon maps. The assumed source parameters are $\Gamma_l = 1.7$, $\Gamma_h = 2.5$, $E_p = 1$ TeV, $\theta_{\text{jet}} = 3^\circ$, $\Theta = 5^\circ$, with on-axis fluence 5000 ph, a background photon density of 50 ph/deg², and a magnetic field oriented 60° to the line of sight and 170° from the horizontal axis.

produce more compact image features pursuant to their smaller deflection angles. Typically, to produce an observable ICC halo feature that extends beyond the source mask. While this condition depends on gamma-ray energy, in practice this requires halos larger than roughly 0.6° , and thus we require

$$B \gtrsim 4 \times 10^{-17} \left(\frac{E_{\text{IC}}}{10 \text{ GeV}} \right) (1+z)^4 \text{ G}. \quad (7)$$

For smaller magnetic fields the ICC halo structure will typically be overwhelmed by the direct emission component and its anisotropy substantially degraded by the *Fermi* PSF. These produce ICC halos that for the *Fermi* blazars are characterized by strong anisotropy with an extent dictated by the magnetic field strength, an example of which is shown in Figure 6. This differs from the previous scenario primarily in the origin of the image structure – here not due to the anisotropy of the gamma-ray emission but rather the anisotropy in the pair distribution function and the strong beaming of the inverse-Compton emission.

As with the previous case we are ultimately interested in producing mock realizations of the *Fermi* sky, which we assume is comprised of a halo, direct emission, and background components. Thus, for the large λ_B case we will require all seven of the parameters in the small λ_B case, as well as parameters describing the IGMF strength and orientation. That is the mock *Fermi* images in the presence of an large-scale IGMF are characterized by ten quantities:

- The source redshift, z .
- The 1 GeV–100 GeV fluence, N_{35} .
- The low-energy photon spectral index, Γ_l .
- The high-energy photon spectral index, Γ_h .

- The energy of the spectral break, E_p .
- The gamma-ray jet opening angle, θ_{jet} .
- The gamma-ray jet viewing angle, Θ .
- The local background photon density, N_B .
- The IGMF, \mathbf{B} .

As before, some of these are obtained from values reported for a sub-sample of *Fermi* AGN while others must be sampled from the appropriate distributions; these are discussed in detail in Section 4. In addition we must define \mathbf{B} . While we will review this in Section 4.4, here we note that we do this by specifying independently an orientation and magnitude with the latter set via the current IGMF strength, B_0 .

3.3. Fermi Point Spread Function

We assume the same PSF as described in Broderick et al. (2016), to which we direct the interested reader for details on implementation, and only summarize salient points here.

Because ICC halos have yet to be unambiguously detected, we consider the Pass 8R2_V6 ULTRACLEANVETO photon sample; these are the photons that are confidently associated with an astronomical origin and not necessarily nearby bright sources. The form of the Pass 8R2_V6 PSF is described in Broderick et al. (2016) and for the events of interest here substantially simplified by the weak PSF dependence on energy above 1 GeV and the fact that the collection of events within the Pass 8R2_V6 ULTRACLEANVETO sample is distributed among a large number of potential bore angles. As a result the collective PSF for the Front and Back detectors are well approximated for each by that at a single bore angle bin, corresponding to 36.9° – 45.6° in both cases (for details see Broderick et al. 2016).

In principle, the square geometry of the LAT imposes a strong dependence on the azimuthal angle of the photon (Ackermann et al. 2012a). However, in practice the long duration of the *Fermi* observations (8 years) combined with the solar tracking and eight-fold symmetry result in a nearly cylindrical symmetry (Ackermann et al. 2012a). This may be broken for short duration or bursty events, and thus if the gamma-ray AGN of interest underwent periods of substantial variability a small residual angular structure may appear. However, as discussed in Appendix D, such structure will enter first at the hexadecapole, i.e., $m = 4$, mode, and therefore is easily distinguishable from that due to ICC halo structure.

3.4. Source-Halo-Background Confusion

The direct emission from the source, background, and ICC halo are not spatially distinct. For large-scale uniform IGMF geometries in particular the halos are strongly centrally concentrated, and therefore will suffer from confusion with the source photons. Therefore, applying the observational constraints provided by the known source and background source counts requires a method to partition the ICC halo component between the source and background in a self-consistent manner.

The LAT PSF provides a natural definition of those events that would be identified as “source” photons. Any substantial emission component beyond the 68% containment radius of the Pass 8R2_V6 ULTRACLEANVETO PSF would be identified as extended and therefore not included in the point-source flux estimates. This is also consistent with the energy- and detector-dependent mask that we apply to the images to reduce Poisson noise (see Section 5.3).

Therefore, we implicitly set the normalization of the ICC halo component by generating image realizations with an appropriate number of “source” photons inside the appropriate 68% containment radius, including those from all components. For strong sources with weak halos this makes little difference. For weak sources with strong halos (e.g., those with very hard VHEGR SEDs) this curtails the halo emission appropriately.

Extending the “source” region farther begins to rapidly increase the angular size of the region as a result of the broad power-law tails on the Pass 8R2_V6 ULTRACLEANVETO PSFs. However, we have verified that extending this to the 95% containment radius makes little difference to our ability to detect ICC halos.

3.5. Near-Source Halo Suppression

A small subset of *Fermi* AGN are closer than D_{pp} , and therefore the assumption that the sources were sufficiently far for the full ICC halo is violated. It is possible to generate ICC halo models in this case for which the region contributing to the ICC halos is restricted to that between the Earth and the source. However, for the large-scale, uniform IGMF models the small-angle contributions to the ICC halos are nearly uniformly distributed along the line of sight, enabling a simpler optical depth correction. That is, we reduce the anticipated halo flux by the energy-dependent factor $1 - e^{-D_p/D_{pp}}$, where D_p is the proper distance to the VHEGR source. For the small-scale, tangled IGMF models this over-reduces the contribution for blazar sources arising from the counter-jet; nevertheless, we conservatively adopt the same optical depth correction factor.

3.6. Time Delays and Duty Cycles

Generally, the contributions to the ICC halos at different positions on the sky are not contemporaneous – there is a delay

between ICC halo gamma rays produced along the line of sight and those off. The typical delay times are geometric in nature and therefore correlated with the angular diameter distance D_A from the central gamma-ray source:

$$\delta t \approx \frac{D_A^2}{2D_{pp}c} \theta^2. \quad (8)$$

where θ is the angular size of observed halo. Therefore, the magnitude of the delay anticipated is limited by size of the ICC halo. For ICC halos from gamma-ray blazars the ICC halo is limited by both the size of the magnetic field deflections and the width of the gamma-ray jet. For the latter this gives $\theta \lesssim \theta_{jet} D_{pp}/D_A$, hence conservatively

$$\delta t \lesssim \frac{D_{pp}}{2c} \theta_{jet}^2. \quad (9)$$

For a typical $\theta_{jet} \approx 3^\circ$ and $D_{pp} \approx 300$ Mpc at $z = 0.3$ for a TeV VHEGR this gives $\delta t \lesssim 10^6$ yr. While considerably larger than the present observing time, this is comfortably short in comparison to the typical radio duty cycles of a few times 10^7 yr to a few times 10^8 yr (Alexander & Leahy 1987; Nulsen et al. 2005; McNamara et al. 2005; Shabala et al. 2008), suggesting that the current gamma-ray flux is indicative of that responsible for a putative ICC halo.

4. THE *Fermi* GAMMA-RAY BLAZAR TARGETS

Here we summarize a variety of essential properties of the observed *Fermi* sample of bright, nearby gamma-ray blazars, i.e., objects likely to have detectable ICC halos. Among these are the intrinsic source parameters, e.g., observed flux, redshift, etc., as well as extrinsic source context, e.g., local gamma-ray background, PSF, and the putative IGMF. Some of these can be specified for each source explicitly, others must be determined probabilistically. The list of targets that met all requirements within the *Fermi* Pass 8R2_V6 ULTRACLEANVETO class is presented in Table 1.

Because ICC halos are essentially reprocessed VHEGR emission, they exist solely around VHEGR-bright objects. Therefore, the goal of initial source class identification is to estimate VHEGR brightness of individual gamma-ray AGN. To do this we exploit the 2FHL, which is optimized for the detection of objects above 50 GeV (Ackermann et al. 2016a). We further restrict our attention to objects that also appear in the 3LAC (Ackermann et al. 2015b) and have a measured redshift, yielding 122 objects. These are dominated by BL Lac-like (BLL) objects, as opposed to the flat-spectrum radio quasars that compose the vast majority of the remainder of the *Fermi* AGN population. This is consistent with the strong correlation between AGN type and spectral hardness, for which reason BLLs typically dominate at high energies.

The *Fermi* PSF provides a direct lower-limit on the size of any detectable ICC halo, both by smearing the anisotropic structure and through contamination from the much brighter direct emission from the source. Beyond $z = 1$ the angular size of the region over which pairs are efficiently produced, i.e., the angular size of D_{pp} , is near the *Fermi* PSF. A more stringent limit comes from the typical deflections in a large-scale IGMF:

$$\Delta\alpha_{def} \approx \frac{\omega_B t_{IC}}{\gamma^2} \approx 1.5 \left(\frac{B}{10^{-16} \text{ G}} \right) \left(\frac{E_{IC}}{10 \text{ GeV}} \right)^{-1} (1+z)^{-4} \quad (10)$$

where $\omega_B \equiv eB/m_e c$ and $t_{IC} \equiv 3m_e c/4\sigma_T u_s \approx 2.4 \times 10^{12} (1 +$

$z)^{-4}$ yr, which sets the size of the ICC halos generated by emission beamed towards us. Noting that the IGMF strength today is reshifted relative to that at high z , for a fixed current IGMF magnitude $\Delta\alpha_{\text{def}} \propto (1+z)^{-2}$ (see Section 4.4). Thus for a current 10^{-16} G IGMF, when $z > 0.6$ the halo size is typically smaller than $0^\circ 6$, and hence likely to be confused with the central source. This is a moderately strong function of the IGMF strength and ICC halo gamma-ray energy, growing to $z \lesssim 4$ for a present-day 10^{-15} G IGMF or 1 GeV ICC halo photons. However, in those cases the halo size is limited by the jet width (Section 4.3). For this reason we impose a limit of $z \lesssim 0.5$, yielding a set of 84 sources.

The above comprise the intrinsic source selection criteria: appearance in the 2FHL and 3LAC with source identification and known redshifts.

4.1. Intrinsic SEDs

The tabulated SEDs in neither the 3FGL (the parent catalog of the 3LAC Acero et al. 2015) nor 2FHL are good estimators of the *intrinsic* TeV brightness for at least two reasons. First, curvature in the SED within the energy bands for which the photon spectral indexes are reported in the 3FGL (100 MeV–100 GeV) and 2FHL (50 GeV–2 TeV) makes any extrapolation to the VHEGR band of interest, 1 TeV–10 TeV, highly uncertain. Second, absorption on the EBL for sources with $z > 0.1$ can have a substantial impact on the SED above 100 GeV, rendering the *observed* VHEGR SEDs poor estimators of the *intrinsic* VHEGR SEDs.

For a handful of sources there exist reported deabsorbed SEDs from air Cherenkov telescopes (e.g., MAGIC, VERITAS, H.E.S.S.). However, these suffer from a number of additional limitations. Typically they provide measurements over a very limited temporal window of the highly-variable emission from gamma-ray bright blazars, and hence are often poor estimators of the time-averaged fluence over long periods. Additionally, the deabsorption prescriptions vary substantially among sources and thus they do not provide a homogeneous class of SED estimates. Finally, for all but the brightest sources the reported Cherenkov telescope SEDs are limited to below the VHEGR energy band of interest, producing the same uncertainties that arise from the 2FHL and 3FGL. They do, however, indicate the degree of variability we may expect over decadal timescales.

Therefore, we independently generate composite SEDs by collating the 3FGL and 2FHL band-specific energy flux measurements for each of the 84 common sources with redshifts below 0.5 as described in Appendix E. That is, gamma-ray SEDs were produced by compiling the 0.1–0.3 GeV, 0.3–1 GeV, 1–3 GeV, 3–10 GeV, and 10–100 GeV flux measurements reported in the 3FGL and 50–171 GeV, 171–585 GeV, and 585–2000 GeV reported in the 2FHL. The former (3FGL) are averaged over 4 years (Aug 2008–Jul 2012) and the latter (2FHL) are averaged over 7 years (Aug 2008–Aug 2015), and thus present nearly contemporaneous ranges extending over many years.

These were deabsorbed using the pair-creation optical depth given by

$$\tau_{\text{obs}}(E, z) = \int_0^z dz' \frac{dD_P/dz}{D_{\text{pp}}[E(1+z), z]}, \quad (11)$$

where D_P is the proper distance to the source, evaluated at the geometric center of the energy bin, i.e. we set

$$F_{E, \text{deabs}} = F_{E, \text{obs}} e^{\tau_{\text{obs}}(E, z)}, \quad (12)$$

where $F_{E, \text{obs}}$ is the specific energy flux. Both the observed and deabsorbed SED for Mkn 421 are shown in Figure 7, with the remainder of the sources used here (i.e., those listed in Table 1) shown in Appendix E.

The deabsorbed SEDs are generally well fit by a broken power-law SED, defined by a normalization, a pivot energy E_p , and photon spectral indexes above and below E_p , Γ_h and Γ_l , respectively:

$$\frac{dN}{dE} = f_0 \begin{cases} (E/E_p)^{-\Gamma_l} & E \leq E_p \\ (E/E_p)^{-\Gamma_h} & E > E_p \end{cases}. \quad (13)$$

A maximum-likelihood fit of the broken-power-law SED model was performed to each candidate source. The result is also shown by the dashed line in Figure 7 for Mkn 421. We have visually verified that small variations in the assumed initial starting point results in negligible variations in the final fit parameter values (though large variations can result in erroneous fit results).

Two classes of qualitatively different SEDs were found from the deabsorbed spectral fits: sources with spectral breaks that are convex ($\Gamma_h > \Gamma_l$) and those that are concave ($\Gamma_h < \Gamma_l$). The former are consistent with the expectation from single-zone inverse-Compton models of VHEGR sources (Ghisellini et al. 1998; Abdo et al. 2010c; Ackermann et al. 2016a). The latter suggest the need for an additional spectral component, either due to an additional comptonizing population or alternative emission source (Böttcher et al. 2013; Cerruti et al. 2015; Zacharias & Wagner 2016).

Uncertainties on the fit parameters were obtained via a Monte Carlo analysis. Trial fluctuations in the fit parameters were constructed from normal distributions with standard deviations given by the Fisher matrix error estimates. An estimate of the allowed range was obtained by taking the collection of parameter values for which the log-likelihood, i.e., χ^2 , increased by unity, shown by the gray bands in Figure 7. This was especially important for sources with only upper limits at high energies, and thus for which the uncertainty in the high-energy SEDs were highly asymmetric.

Generally, the normalization, E_p , and Γ_l were tightly constrained; this is in part a selection effect as each object is a well-characterized *Fermi* source. Thus in our set of bright, nearby gamma-ray bright AGN we fix these to their observed (normalization) or fitted (E_p and Γ_l) values. In contrast, Γ_h is considerably more uncertain, often as a result of a high E_p and larger uncertainties or upper limits on the intrinsic high-energy flux estimates. Therefore, for the purpose of generating ICC halo realizations we stochastically choose Γ_h over the permitted range. Because this is typically asymmetric, permitting either much smaller or larger values of Γ_h , we assume that the probability of a given Γ_h is well approximated by two one-sided normal distributions centered at the best fit value with standard deviations set by the range obtained by the Monte Carlo procedure above and below.

Key intrinsic target parameters that enter the generation of mock realizations of the sample are the number of observed source photons, z , and the source intrinsic SED fit parameters. These are listed in Table 1 for the *Fermi* targets that are used in this paper (see also Section 5.6).

4.2. Local Gamma-ray Neighborhood

In addition to the intrinsic source requirements, the diffuse, large-scale nature of the ICC halos places constraints on the neighborhood of targets. While these are essentially limits on




Figure 7. Composite SED for Mkn 421 taken from the 3FGL (red) and 2FHL (dark blue) catalogs and their deabsorbed counterparts (orange and light blue, respectively, though the former are indistinguishable from the observed values). A broken power-law fit is shown by the dashed line, the grey band denotes the uncertainty in the SED reconstruction. For comparison VHEGR observations by ground-based air Cherenkov telescopes (Whipple observations of the high state in 2000/2001, H.E.S.S. observations in 2004, and MAGIC observations in 2004/2005, top to bottom, adapted from Figure 10 of Albert et al. 2007) are shown in purple (observed) and magenta (deabsorbed).

potential contaminating features, for the purpose of target selection this reduces neighboring sources and large-scale background gradients.

Even weak neighboring sources can produce a large bias in the angular power spectra. While we defer a characterization of this signal to Section 5.4.2, an initial target-list cut was made to remove all sources with bright neighbors within 2° . Beyond 2° neighbors were permitted, though the area over which the power spectrum analysis was performed was restricted to prevent contamination. This is illustrated in Figure 8, which shows examples of excluded, restricted, and ideal sources.

The gamma-ray background varies substantially from source to source as a result of the different sky location. This is dominated by the diffuse Galactic component, and becomes noticeably worse at low Galactic latitudes, where it imparts substantial gradients in the gamma-ray counts. Rather than attempting to model this component we exclude sources with strong background gradients visible over scales of 10° (typically corresponding to $|l| < 18^\circ$) and make an estimate of the background photon density for each target source individually. In practice, the presence of a background gradient appears dominantly in the dipolar power, and thus is distinguishable from the bipolar signals of interest (see Section 5.4.2).

Therefore, the above comprise the extrinsic source selection criteria: no neighbors within 2° , restricting from 4° as necessary, and no large-scale gradients in the background flux. Key extrinsic target parameters that enter the generation of mock realizations of the sample are the maximum non-contaminated angular radius, R_{\max} , and the total number of front- and back-converted photons within the permitted region due to the source and background. There are 27 objects that are sufficiently isolated and satisfy the intrinsic source parameter selection criteria, which are shown in Tables 2 and 3. Final selection of the *Fermi* targets listed in Table 1 will be described in section 5.6.

4.3. Jet Opening and Viewing Angles

For most *Fermi* blazars the viewing and intrinsic opening angles are poorly constrained. However, the subset of objects that also appear in the sample monitored by the MOJAVE group provides some guidance on the parsec-scale radio opening and viewing angles (Pushkarev et al. 2009). Generally, *Fermi* blazars exhibit radio jets that are intrinsically more narrow than the gamma-ray dim population. Figure 9, adapted from Pushkarev et al. (2009), shows the distribution of intrinsic radio opening angles (full-width half-max, FWHM), which peaks near the median of $1^\circ 0$ and falls rapidly thereafter. This is well-fit by a generalized λ distribution with parameters given in Appendix A, shown in the figure by the red line.

Interpreting the implications for the gamma-ray jet opening angles is complicated for a number of reasons. First, obtaining the radio opening angles is itself challenging. While measuring the apparent (projected) opening angle is straightforward, determining the intrinsic opening angle typically requires kinematic information obtained from multiple widely-separated epochs of imaging observations. Second, the radio and gamma-ray opening angles need not be the same, and are generally different. Moreover, the current uncertainty in the gamma-ray emission mechanism precludes using spectral information to relate the two directly.

To assess the relationship between the two we begin with the following assumptions:

1. The flux distribution of *Fermi* sources is limited from above, i.e., they exhibit a maximum intrinsic luminosity.
2. The intrinsic opening angle of the gamma-ray and radio jets are proportional, i.e., broad radio jets are broad gamma-ray jets and vice versa.
3. The gamma-ray jet is structured as a Gaussian with a source-dependent standard deviation, θ_{jet} . That is, the gamma-ray flux within the jet observed at an angle Θ is given by

$$\mathcal{F} = \mathcal{F}_0 e^{-\Theta^2/2\theta_{\text{jet}}^2}. \quad (14)$$

The first is a good approximation in practice since *Fermi* only sees the bright end of the hard gamma-ray blazar luminosity function. The second is natural given that the overwhelming majority of gamma-ray bright AGN are blazars (Ackermann et al. 2011, 2015b). The quantitative consequence is that the gamma-ray jet scale, θ_{jet} , is related to the intrinsic radio jet FWHM, α_{int} , by some proportionality constant:

$$\theta_{\text{jet}} = \sigma_{\text{jet}} \alpha_{\text{int}}. \quad (15)$$

The third is already made in the gamma-ray halo models described in Broderick et al. (2016) and employed here. Importantly, note that the assumed structure is effectively a condition on the *apparent* jet opening angle; while generally $\tan \alpha_{\text{app}} = \tan \alpha_{\text{int}} / \sin \Theta$, for small radio opening and viewing angles this is approximately $\tan \alpha_{\text{app}} \approx \alpha_{\text{int}} / \Theta$, and hence the condition imposed by the third can be cast in terms of the *apparent* opening angle alone:

$$\mathcal{F} = \mathcal{F}_0 e^{-1/2 \sigma_{\text{jet}}^2 \tan^2 \alpha_{\text{app}}}. \quad (16)$$

This has two consequences. First, the number of *Fermi* blazars for which α_{app} is known is comparatively large. Second, because α_{app} is directly measured during a single epoch of radio observations it is much better known than α_{int} or Θ .




Figure 8. Fluence maps of example excluded (due to neighboring sources, left), restricted (center), and ideal (right) *Fermi* sources. Color indicates the number of photons within a given pixel. Dashed circles have radii of 2° , 3° , and 4° .




Figure 9. Intrinsic parsec-scale radio opening angle (FWHM) distribution for *Fermi* sources in the MOJAVE data base (adapted from Pushkarev et al. 2009). Shown by the red line is the fitted generalized Λ distribution from which we draw intrinsic parsec-scale radio opening angles.




Figure 10. Flux vs. viewing angle in units of radio jet width, approximated by $(\tan \alpha_{\text{app}})^{-1}$ where α_{app} is the apparent parsec-scale radio opening angle for hard ($\Gamma_l < 2.1$, blue) and soft ($\Gamma_l > 2.1$, red) *Fermi* sources in the MOJAVE data base. Fits to the upper-envelope assuming a Gaussian jet are shown by the light-blue filled and light-red filled regions.

While the value of \mathcal{F}_0 is unknown for any given source, the presence of an upper limit implies that *Fermi* sources should populate the region in the \mathcal{F} - α_{app}^{-1} plane defined by some $\mathcal{F}_{0,\text{max}}$. Show in Figure 10, this is the clearly the case – sources with large α_{app}^{-1} have systematically lower fluxes, falling in a manner consistent with the Gaussian dependence posited. This presents a direct way in which to measure σ_{jet} by fitting the envelope of points in the \mathcal{F} - α_{app}^{-1} plane.

Interestingly, the value obtained depends strongly on spectral hardness. Softer sources ($\Gamma_l > 2.1$) have systematically broader gamma-jets, with $\sigma_{\text{jet}} \approx 2.5$; the hard sources of interest here ($\Gamma_h < 2.1$) have 45% narrower jets, with $\sigma_{\text{jet}} \approx 1.6$. It is this latter relationship we adopt.

Note that this implies that the gamma-ray emission is beamed over a substantially larger angle than subtended by the parsec-scale radio jet. The ratio of the FWHMs of the gamma-ray and parsec-scale radio jets is $(8 \ln 2)^{1/2} \sigma_{\text{jet}} = 3.77$, diffusing the emission over a solid-angle nearly 14 times larger. This suggests that the beaming of the gamma-ray emission is similar, qualitatively, to the presumed jet structure near its base, i.e., within the collimating region.

This does not mean, however, that we anticipate large beaming corrections to the apparent gamma-ray flux or that there should be a large population of non-blazar gamma-ray sources observed by *Fermi*. Generally, parsec-scale jets are more tightly collimated in comparison to the typical beaming angle of the radio emission, with

$$\alpha_{\text{int}} \approx \frac{0.3}{\Gamma_{\text{jet}}}, \quad (17)$$

where Γ_{jet} is the jet Lorentz factor (not to be confused with a photon spectral index); this continues to hold for the *Fermi* subset of MOJAVE sources (Pushkarev et al. 2009). That is, the typical angular scale over which the radio emission is beamed, Γ_{jet}^{-1} , is itself three times the parsec-scale jet FWHM. Comparing the gamma-ray and radio beaming angle gives $\theta_{\text{jet}} \Gamma_{\text{jet}} \approx 0.48$ for the hard gamma-ray blazars. That is, for these sources the gamma-ray emission is moderately more beamed than the radio emission. Despite having a $\theta_{\text{jet}} \Gamma_{\text{jet}}$ that is roughly 40% larger, this remains true for the soft gamma-ray blazars – the gamma-ray emission is more beamed than the radio emission.

Because we do not have an a priori estimate of either θ_{jet} or Θ for any particular source, for the purpose of creating ICC halo realizations we randomly generate values using the following procedure:

1. We select a value of α_{int} from the distribution shown in Figure 9 and described in Appendix A.
2. Based on α_{int} we estimate the relativistic beaming angle to be $\Gamma_{\text{jet}}^{-1} \approx \alpha_{\text{int}}/0.3$. To appear as a blazar we must be viewing the source within this angle, and thus choose Θ from an isotropic distribution within Γ_{jet}^{-1} .
3. We generate $\theta_{\text{jet}} = \sigma_{\text{jet}}\alpha_{\text{int}}$ with $\sigma_{\text{jet}} = 1.6$.

4.4. IGMF Strength and Orientation

In the presence of a large-scale IGMF the ICC halo images depend upon the assumed strength and orientation of the IGMF. Without any prior knowledge of either the strength or orientation of the IGMF we consider variations in both. In the case of the former we set the strength of the IGMF today, B_0 , and assign the local strength at the source to be the appropriately redshifted value, i.e., the local magnetic field strength is

$$B = B_0(1+z)^2. \quad (18)$$

For a given analysis we assume that all sources have the same value of B_0 , assessing the detectability of the ICC halos as a function of field strength. Because we restrict our attention to $z < 0.5$ this has at most a factor of two impact on the assumed B .

Realizations for the orientation of the field is generated from an isotropic distribution for each image independently. That is, while in this limit we assert that the field is coherent over scales large in comparison with the gamma-ray jet widths, we permit large variations over the distances between *Fermi* AGN in our sample. This is consistent with a picture in which the ICC halos are produced in cosmological voids in which the IGMF has been imprinted from early times. We note that this is rather pessimistic – correlations in the orientations of ICC halos from neighboring VHEGR sources would permit coherent stacking, which we ignore here.

5. ASYMMETRIC SIGNATURES OF HALOS

As introduced in Section 2 we exploit the anticipated structure in the ICC halos by constructing a statistical measure of the anisotropy. To do this we focus on the angular power spectra defined in terms of an event-specific polar angle, θ_j , about the source center, given in Equation 5. This is a natural statistic for the reasons discussed in Section 2: it is sensitive in particular to anisotropic structure yet independent of an absolute rotation of the source orientation. Moreover, it is only weakly contaminated by the known systematics in the *Fermi* LAT instrument responses and the background source population, and in a way that is easily distinguishable from the ICC halo signal of interest. There are, however, a number of practical steps, which we discuss here. These include removing coordinate aberration from the *Fermi* images about a source (Section 5.1), identifying the source center robustly (Section 5.2), and masking source contamination (Section 5.3). We then stack the power spectra from multiple sources (Section 5.5) and determine our final optimized source list (Section 5.6).

5.1. Converting to Locally Euclidean Coordinates

At large latitudes coordinate aberration induces angular structure in the images of even cylindrically symmetric sources. Therefore, we must first approximately flatten the gamma-ray images from *Fermi*, i.e., transform to a set of flat

coordinates. Because we are interested only in the angular structure of the gamma-ray distribution about the central source (α_s, δ_s), i.e., we are not concerned with its radial structure, fully flattening the images is unnecessary. Rather it is sufficient to remove the angular distortion.

To do this we begin with the gamma-ray positions in equatorial coordinates (α_j, δ_j) and perform a rotation to align the reported source position along the polar axis. That is on the unit sphere we set the gamma-ray positions to

$$\begin{aligned} x_j &= \sin \delta_s \cos \delta_j \cos(\alpha_j - \alpha_s) - \cos \delta_s \sin \delta_j \\ y_j &= \cos \delta_j \sin(\alpha_j - \alpha_s) \\ z_j &= \cos \delta_s \cos \delta_j \cos(\alpha_j - \alpha_s) + \sin \delta_s \sin \delta_j, \end{aligned} \quad (19)$$

from which we obtain the angular positions by projecting along \hat{z} :

$$\alpha'_j = \frac{x_j \cos^{-1}(z_j)}{\sqrt{x_j^2 + y_j^2}} \quad \text{and} \quad \delta'_j = \frac{y_j \cos^{-1}(z_j)}{\sqrt{x_j^2 + y_j^2}}. \quad (20)$$

5.2. Maximum-Likelihood Center Finding

The source positions reported by *Fermi* provide an excellent initial estimate of the source locations. However, these are obtained from a different set of event reconstructions (Pass 7R_V15 SOURCE) from those employed here (Pass 8R2_V6 ULTRACLEANVETO). Moreover, even potentially small offsets provide an obvious systematic uncertainty that will generate spurious power at $m = 1$. While we discuss how such an error may be naturally identified in the structure of the angular power spectra directly in Section 5.4.2, we also make an effort to mitigate this directly via an improved estimate of the source location. This also has the effect of treating the mock images in the same fashion as the real data – both will have small offsets in the source location based on the particular photon realization.

To identify the source location we make a maximum-likelihood estimate of the source location, (α'_s, δ'_s) . The likelihood was taken to be composed of a Gaussian source of known size on top of a uniform background within a specified angular size on the sky. The result is an estimate for the source location and ratio of the source-background fluences. Additional details of the method can be found in Appendix B.

While the *Fermi* PSF is neither Gaussian nor independent of energy, and thus our likelihood does not formally describe the *Fermi* response for a point source, both have proven to be adequate approximations for our purpose, producing highly accurate source location estimates. We verified this by generating mock point source images following the algorithm described in Section 3, which utilizes the fully energy-dependent *Fermi* Pass 8R2_V6 PSF, and generating source location estimates. The distribution of offsets is compared to the *Fermi* Pass 8R2_V6 PSFs at low and high energies in Figure 11. In all cases our estimate is roughly an order of magnitude better than the characteristic width of the *Fermi* Pass 8R2_V6, a reflection of the large number of photons in the gamma-ray maps of the bright sources of interest.

We generate a final set of source-centered positions

$$\alpha''_j = \alpha'_j - \alpha'_s \quad \text{and} \quad \delta''_j = \delta'_j - \delta'_s, \quad (21)$$

which we convert into polar coordinates:

$$r_j \equiv \sqrt{\alpha''_j{}^2 + \delta''_j{}^2} \quad \text{and} \quad \theta_j = \tan^{-1}(\delta''_j / \alpha''_j). \quad (22)$$




Figure 11. Distribution of displacements between the true and reconstructed centers for a number of mock sources with typical fluences in the 1 GeV–100 GeV band. For reference the Pass 8R2_V6 ULTRACLEANVETO PSFs for front-converted events at 1 GeV (red) and 100 GeV (blue) are shown. In all cases the reconstructed center locations are at least an order of magnitude better than the typical PSF width.

It is these θ_j that enter into Equation (5) to construct the object-specific angular power spectra.

5.3. Source Masking and Contamination Mitigation

The primary sources of noise in the power spectra estimates are the photons from the source and in the background. Beyond careful source selection, e.g., removing objects with obvious contaminating neighbors or strong background gradients, there is little that can be done regarding the background. However, this is not true for the source itself.

Any structure on scales comparable to the PSF width will be erased, eliminating the value of photons near the source. At the same time, the direct photons from a point source contribute dominantly within this region. Therefore, prior to computing the angular power spectrum for each object we apply an energy dependent mask, excising the region inside the 68% containment radius of the Pass 8R2_V6 ULTRACLEANVETO PSF at each energy, independently for the front and back detectors. This eliminates at once both a region without a significant anisotropy signal and a substantial source of noise. We assess the impact of variations in the size of the excluded region in Section 6.3, generally finding that it is negligible.

For most objects the ICC halos extend over angular scales that are large in comparison to the mask, rendering the mask moot. However, for present-day IGMF strengths less than $3 \times 10^{-17}(1+z)^2$ G, corresponding to either weak fields or high- z , the ICC halo can lie completely within the masked region.⁴ In effect this simply extends the constraint on ICC halo size already imposed by the finite resolution of the LAT marginally.

5.4. Example Single-Source Power Spectra

Example average power spectra for a single, bright source are shown in Figures 12–15. These include both the power

⁴ This limit arises from combining the angular scale implied by a deflection in Equation (10) and the cosmological evolution of a fixed strength field, given in Equation (18) assuming that the source is sufficiently close that the angular diameter distance is similar to D_{pp} . Thus, we set $\alpha_{\text{def}}/2 \approx 0.2$, comparable to the typical mask size near 10 GeV, from which we obtain $B_0 = B(1+z)^{-2} = 3 \times 10^{-17}(1+z)^2$. Note that the *observed* IC gamma-ray energy in Equation (10), E_{IC} , does not redshift.




Figure 12. Median mock power spectra for a bright *Fermi* source containing a large-scale, uniform IGMF constructed from front-converted events (red +’s), back-converted events (green x’s), and all events (blue circles). The assumed IGMF strength is 10^{-15} G and $\Gamma_l = 1.7$, $z = 0.3$, $\Theta = 5^\circ$, $\theta_{\text{jet}} = 3^\circ$, $\Gamma_h = 2.5$, $E_p = 1$ TeV, with on-axis fluence of 5000 ph, a background photon density of 50 ph/deg², and a magnetic field oriented 60° to the line of sight and 170° from the horizontal axis. At large m all cases asymptote to the Poisson limit, $P_m \approx N^{-1}$.




Figure 13. Median mock power spectra for a bright *Fermi* source assuming different IGMF structures. Power spectra for a large-scale, uniform IGMF with strength 10^{-15} G (red squares), 10^{-16} G (blue circles), 10^{-17} G (green crosses) are shown; decreasing field strength results in a smaller low- m power excess. In all cases $\Gamma_l = 1.7$, $z = 0.3$, $\Theta = 5^\circ$, $\theta_{\text{jet}} = 3^\circ$, $\Gamma_h = 2.5$, $E_p = 1$ TeV, with on-axis fluence 5000 ph, a background photon density of 50 ph/deg², and a magnetic field oriented 60° to the line of sight and 170° from the horizontal axis.

spectra anticipated from the various ICC halo models under consideration and those from a variety of potential contaminants. Analytical computations for approximate cases of potential relevance are also collected in Appendix C.

5.4.1. ICC Halos

Power spectra may be made for events converted separately in the front, back, or entirety of the LAT. While front-converted events have a smaller PSF, and therefore maintain small-scale power more effectively, the improvement in event statistics arising from the near-doubling of the gamma-ray number when back-converted events are included produces an overall improvement in the ability to identify ICC halos. This is clearly evident in the comparison between the power at small and large




Figure 14. Median mock power spectra for a bright *Fermi* source assuming a small-scale, tangled IGMF with $B_0 \gtrsim 10^{-14}$ G for objects seen at large ($\Theta = 60^\circ$, red squares) and small ($\Theta = 5^\circ$, blue circles) viewing angles. The power spectrum in the absence of an ICC halo is shown for comparison in green crosses. In all cases $\Gamma_l = 1.7$, $z = 0.3$, $\theta_{\text{jet}} = 3^\circ$, $\Gamma_h = 2.5$, $E_p = 1$ TeV, with on-axis fluence 5000 ph and a background photon density of 50 ph/deg².

m in Figure 12, which shows the each power spectrum class individually.

In all cases at large m the \mathcal{P}_m reach the Poisson noise limit, producing a characteristic flattening at $\mathcal{P}_m \approx N^{-1}$, indicating the effective number of image photons used. The key discriminant that provides evidence for ICC halos is the disparity between \mathcal{P}_2 and this floor, for which the power spectrum constructed from all events is largest. Thus, henceforth, we show only the power spectra for the entire event list, i.e., including both front- and back-converted events.

A comparison of the ICC halo signal for different assumptions regarding the IGMF and orientation are shown in Figures 13 and 14. In contrast to the null case, the clear signal for an ICC halo is the large quadrupole power, i.e., \mathcal{P}_2 , in comparison to the Poisson limit. Moreover, the clear oscillatory nature is a signature for the anticipated near bipolar symmetry. Importantly, this “sawtooth” structure is a key systematic diagnostic, differentiating a true ICC halo signal from potential power spectrum contaminants (see the following subsection).

Nevertheless, the magnitude and structure of the power spectrum is strongly sensitive to the assumed IGMF geometry and source viewing angle. For acute viewing angles, i.e., $\Theta \approx \theta_{\text{jet}}$, the large-scale, uniform IGMF models are most significantly distinct from the null case. In contrast, in this case the small-scale, tangled IGMF is difficult to distinguish as a result of the extreme foreshortening and dilution. However, for oblique viewing angles, i.e., $\Theta \gg \theta_{\text{jet}}$, the small-scale, tangled IGMF models do exhibit noticeable structure in the power spectrum; we will address possibilities for exploiting gamma-ray observations of oblique jets elsewhere.

For the large-scale uniform IGMF models the power excess extends beyond the even multipoles as a consequence of the breaking of the bipolar symmetry by the intrinsic structure of the jet. When viewed along the jet axis a near-perfect symmetry exists as the positrons and electrons gyrate in opposite directions, generating halo emission on opposite sides of the source. However, when viewed at angles comparable to the jet width one component is suppressed by the comparative deficit in pair production due to the reduced VHEGR flux. This is insufficient to remove the characteristic “sawtooth” pattern in

any case.

Weaker IGMFs produce smaller ICC halo signals, vanishing below 10^{-17} G as a result of the source mask. IGMFs that are much stronger than 10^{-14} G also produce weaker ICC halo signatures in the power spectra as a result of the dilution of the gamma-ray flux due to the multiple gyrations. Thus, in principle, apart from simple detection, which is the focus of this work, it should be possible to characterize a large-scale IGMF given a measurement of the gamma-ray angular power spectra.

Also visible in Figure 13 is a two-zone analog of a gamma-ray excess at large angular radii from the central source, i.e., the signal described in Figure 3. In the stacked angular power spectrum this takes the form of a systematic drop in $\mathcal{P}_m^{\text{tot}}$ at large m that systematically grows with increasing B_0 . This arises from the ICC halos moving photons from within the central masked region to outside, where they are included in the angular power spectrum estimate, decreasing the Poisson noise limit. As discussed earlier, exploiting this signal requires knowing the gamma ray background and radial structure of the PSF to high accuracy a priori.

5.4.2. Contaminants

Low-multipole angular power can also be produced by features that are independent of ICC halos. These include systematic errors in the generation of the angular power spectra, angular structure in the *Fermi* Pass 8R2_V6 ULTRACLEAN-VETO PSF, and unresolved features in the background. Here we quantitatively consider each of these. Importantly, none produce the quadrupole-dominated, sawtooth structure in the angular power spectrum characteristic of the bipolar structures associated with the ICC halos.

As has already been described in Section 5.2, the location of the central source may be identified with an accuracy that significantly exceeds the typical PSF width of *Fermi*. Nevertheless, centering errors combined with the large radial gradients in the gamma-ray flux away from the source center produce a natural source of error in the angular power spectrum (see, e.g., Appendix C.2). To assess the magnitude of this error, we intentionally offset a bright gamma-ray source typical of the *Fermi* sample by $0^\circ 05'$, corresponding to roughly 10% of the typical width of the 1 GeV Pass 8R2_V6 ULTRACLEANVETO-front PSF (top). As seen in the top panel of Figure 15 this leads to spurious power at low multipoles, dominated by the dipole, falling rapidly with m , and joining the typical Poisson noise by $m = 2$. The large dipole-quadrupole power ratio makes this easily distinguishable from a signal attributable to ICC halos.

Unresolved sources within the background surrounding the primary *Fermi* AGN also provide a natural source for dipole angular power. This arises in two instances, the first of which is a single neighboring object just below our exclusion threshold. Given that the *Fermi* detection threshold is 5 ph (Ackermann et al. 2013b), this is unlikely to produce a substantial contribution for images comprised of many thousands of events. Adding a source that is obviously visible (150 ph, intermediate to the companions in the left and center panels of Figure 8) produces a notable excess of power at low multipoles (top-right panel of Figure 15). The degree of this excess depends on the location of the contaminating neighbor, becoming larger for more distant unresolved sources. In practice, this is many times over the *Fermi* detection threshold and would already be excluded. Much weaker companions would produce a correspondingly weaker contribution to the power spectrum. In all cases, as before it is dominated by the dipole component and fails to exhibit the sawtooth morphology.

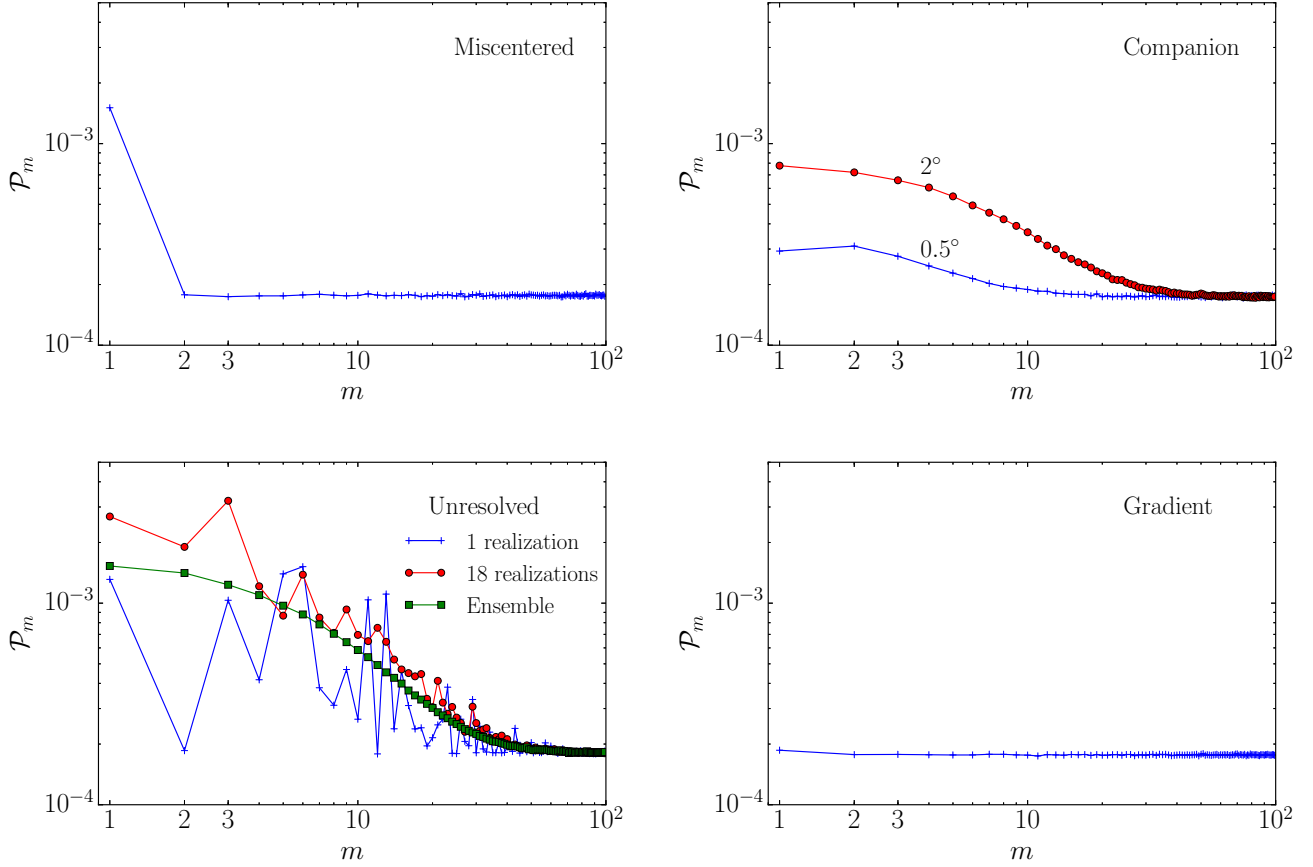


Figure 15. Median mock power spectra for a bright *Fermi* source without an ICC halo but various contaminating systematics. Top left: a miscentered image, shifted by $0^{\circ}05$, corresponding to roughly 10% of the typical width of the 1 GeV Pass 8R2_V6 ULTRACLEANVETO-front PSF. Top right: a faint companion (150 ph) that lies well above the *Fermi* detection threshold (roughly 5 ph) but below the visual inspection threshold, located at $0^{\circ}5$ (blue) and 2° (red). Bottom left: a background comprised of many unresolved gamma-ray point sources consistent with the flux distribution of *Fermi* point sources for a single realization (blue), 18 realizations (red), and an ensemble of realizations (green). Bottom right: a weak gradient in the background comparable to that of the visual inspection threshold. In all cases $\Gamma_l = 1.7$, $z = 0.3$, $\Theta = 5^{\circ}$, $\theta_{\text{jet}} = 3^{\circ}$, $\Gamma_h = 2.5$, $E_p = 1$ TeV, with on-axis fluence 5000 ph and a background photon density of 50 ph/deg 2 .

The second way in which unresolved sources enter is through the origin of the background itself. At high energies (> 50 GeV) nearly the entirety of the extragalactic gamma-ray background has been resolved into point sources (Ackermann et al. 2016b). Extending this to lower energies results in a clustering of photons about the brightest background objects, and therefore low-multiple power arising from a handful of unresolved neighbors⁵. In practice, the extended fields about individual targets in Table 1 do include known 3FGL sources with up to 40 ph, well above the *Fermi* detection threshold. Thus, shown in the bottom-left panel of Figure 15 is the angular power spectrum arising from a background comprised of a population of sources below the 40 ph threshold and distributed according to the 2FGL flux distribution (Abdo et al. 2010a; Broderick et al. 2014b). Because the latter is formally divergent, we filled in the background from the high-flux end, beginning just below the *Fermi* threshold and stopping when the required number of background photons were obtained. For all of the background sources we adopted a photon spectral index of 2.4, consistent with the 1 GeV–100 GeV background (Abdo et al. 2010b; Ackermann et al. 2015a).

⁵ This is extremely pessimistic. Even at high latitudes the smooth, Galactic contribution to the gamma-ray background is substantial. Nevertheless, this gives an extreme estimate of the impact of a highly structured background.

Here it is useful to distinguish between random realizations of the background sources, i.e., the positions and fluxes of the unresolved sources that comprise the background, and random realizations of photons drawn from the sources. Even after averaging over a full ensemble of photon realizations a single background source realization produces a stochastic angular power spectrum exhibiting large variations at low multipoles, shown by the blue line in the bottom-left panel of Figure 15. This is a result of the structure imposed on the background via the locations and strengths of the background sources and cannot be overcome by collecting additional observations. Averaging over realizations of the background sources and photons results in a smooth angular power spectrum, similar to that arising from a nearby, unresolved source (green line in the bottom-left panel of Figure 15).

In principle, we may directly measure this background angular power spectrum using nearby empty fields. However, our ability to remove it is fundamentally limited by the variance due to moderate number of background realizations presented by the source list in Table 1. Nonetheless, even after averaging over 18 sources (the number of sources in Table 1) the fluctuations are already substantially reduced (red line in bottom-left panel of Figure 15). Regardless, the stochastic structure lacks the telltale sawtooth morphology of the ICC halos.

Azimuthal structure of the *Fermi* Pass 8R2_V6 ULTRACLEANVETO PSF will also generate angular power within the source photons. Note that this is not true for the distribution of the uniform gamma-ray background – even a highly anisotropic PSF cannot impart structure on a uniform field of photons. The angular structure of the *Fermi* Pass 8R2_V6 ULTRACLEANVETO PSF arises from the square geometry of the LAT (see also Section 3.3). The instantaneous PSF exhibits only 5% variations in the PSF at the energies of interest (Ackermann et al. 2012a). Over timescales short in comparison to years this is substantially reduced by the rotation of the *Fermi* during solar tracking and the eight-fold symmetry of the LAT.

In Appendix D we make an estimate of the residual PSF-induced angular power, assuming that cumulative gamma-ray image is comprised of many epochs during which the roll angle of *Fermi* is highly correlated. The duration of these epochs correspond to roughly the time for *Fermi* to rotate by 45° , after which it effectively rotates through the entirety of the PSF as a result of the LAT’s square geometry. Presuming that the fluence is relatively evenly distributed over the past 8 yr, even for optimistic assumptions regarding the structure of the PSF, the estimated residual angular power is less than 1% of the anticipated Poisson noise. Moreover, it vanishes identically for $m = 2$ as a result of the LAT’s symmetry, exhibiting power only for $m = 4$ and its harmonics.

Finally, we considered a linear gradient in the background photon density. Because we explicitly select sources for which there is no apparent large-scale gradients in the background, we again set the value to our effective detection threshold, corresponding to a variation in the photon density of 20% across the 4° image. As shown in figure 15, the impact on the angular power spectrum is very small, weakly modifying the dipole power primarily.

It is important to note in summary two key results regarding all of the potential the systematic uncertainties arising from contaminants: First, their *shape* is qualitatively different from the distinct signatures of the bimodal ICC halos. Second, their *magnitude* is far smaller than that expected from the ICC halos arising from a large-scale, uniform IGMF. As a result they should be readily distinguishable.

5.5. Combining Multiple Sources

Finally, to increase the fidelity of the angular power spectrum we stack the estimates from multiple sources. Unlike stacking the images directly, this preserves the anisotropic signal; a rotation of any image, corresponding to setting $\theta_j \rightarrow \theta_j + \varphi$, leaves the \mathcal{P}_m unchanged as may be verified by inspection of Equation (5). It does, however, reduce the intrinsic scatter in the power spectrum estimate.

In principle, this may be optimized via the weighting assigned to individual sources – images with higher numbers of intrinsic photons will produce better intrinsic power spectrum estimates and thus may be given additional weight in the stacking processes. The natural way to do this is the variance-weighting, giving the smallest variance in \mathcal{P}_m at each m . In practice, we found that the dominance of the source counts by a handful of objects (e.g., Mkn 421) led to an associated dominance of the power spectra estimate, eliminating much of the power of the stacking process.

It is also possible to exploit the spectra of the ICC halos and/or the redshift-dependence of the ICC halo extent to provide more optimal weightings. We found that no such effort made a substantial impact on the ability to distinguish the

power in the quadrupole and the neighboring odd multipoles, the key observable for our bimodal ICC halos. The reason is simply that while the ICC halos are typically harder than the gamma-ray background they are only marginally so, limiting the ability to spectrally separate the two components.

Therefore, we define our stacked power spectrum by the arithmetic average of the individual source power spectra:

$$\mathcal{P}_m^{\text{tot}} = \frac{1}{N_{\text{src}}} \sum_{\text{src}} \mathcal{P}_m^{\text{src}}, \quad (23)$$

where $\mathcal{P}_m^{\text{src}}$ is the single-source power spectrum defined by Equation (5).

5.6. Source List Optimization

The ability to generate simulated realizations of the *Fermi* sky enables us to theoretically optimize the list of *Fermi* sources that are ultimately stacked. That is, apart from gross properties of the sources (e.g., SEDs and redshifts), we can select the group most likely to collectively produce the apparent signatures of ICC halos in the stacked angular power spectrum *without* looking at the actual structures of these images.

The SEDs described in Section 4.1 provide guidance on which sources are likely to be bright above a TeV. However, while this is a necessary condition, it is not sufficient to produce bright ICC halos. The halo itself is impacted by the source distance (among other parameters that are marginalized over). The ability to detect the halo is impacted by the local background. Thus, armed with the ability to simulate halos from the 27 sources that are sufficiently isolated, appear in the 2FHL and 3LAC and have known redshifts, we applied a final optimization step designed to maximize the ability to detect ICC halos.

We begin by defining a halo-model specific detection likelihood statistic:

$$P_{\text{det}} = \int dp P_{>\mathcal{P}_2}^{\text{h}}(p) P_{\mathcal{P}_2}^{\text{n}}(p), \quad (24)$$

where $P_{\mathcal{P}_2}^{\text{n}}(p)$ is the probability of finding $\mathcal{P}_2^{\text{tot}} = p$ in the null case, i.e., without a halo, and $P_{>\mathcal{P}_2}^{\text{h}}(p)$ is the cumulative probability associated with $\mathcal{P}_2^{\text{tot}} \geq p$ when a halo is present. This is the probability that the $m = 2$ power from the given halo model exceeds that from the null case, marginalized over the probabilities of both. For the null case, $P_{\text{det}} = 1/2$, i.e., in the absence of an ICC halo the probability that \mathcal{P}_2 will exceed that from the null case is simply 50%. This translates directly into the probability that a given halo model will produce excess power in the quadrupole, i.e., the probability that an ICC halo is detectable. From a collection of N_{real} realizations of the full sample of isolated, hard *Fermi* AGN we estimate P_{det} via

$$P_{\text{det}} \approx \frac{1}{N_{\text{real}}} \sum_{\text{real}} \frac{1}{51} \sum_{m=50}^{100} \Theta(\mathcal{P}_2^{\text{tot}} - \mathcal{P}_m^{\text{tot}}), \quad (25)$$

where $\Theta(x)$ is the Heaviside function and we have used the high- m power spectra as a proxy for the null case. This is necessarily a function of the particular set of sources included, changing as a result of all of the systematic inputs into the construction of the ICC halo models and the stacked power spectra. Thus, for a given halo model we can optimize the list of input sources by maximizing the simulated values of P_{det} .

This maximization is achieved in a restricted sense in practice: we order the list of sources by intrinsic TeV brightness

and then construct a sequence of source lists, each comprised of the preceding list and the next-brightest TeV source. For example, the first list includes Mkn 421, the second, Mkn 421 and 3FGL J2347.0+5142, etc. The value of each source is then evaluated when it is first included; if it increases P_{det} it is kept, if it decreases P_{det} it is removed. This procedure results in a set of optimal sources for each halo model.

This list varies between halo models due to the differences in the ICC halos produced. As a result in Table 1 we present four separate samples corresponding to four different choices of B_0 for the uniform IGMF halo models: $B_0 = 10^{-17}$ G, 10^{-16} G, 10^{-15} G, and 10^{-14} G. In practice they are nested: the first sample is comprised of the brightest four sources, the second sample includes the first and the next six sources, etc. This is a result of the smaller apparent sizes of the halos associated with weaker IGMFs, and thus preferring a brighter and nearer AGN sample. It is this collection of source samples that we will use henceforth.

We performed a similar procedure to optimize the source list for the tangled magnetic fields. This was less successful due to the low intrinsic signal in the power spectra. We found little difference among the various samples chosen and therefore simply adopted the set from the uniform IGMF model with $B_0 = 10^{-15}$ G.

It is important to note that this optimization procedure was performed *entirely* using the *simulated* images. That is, apart from the source SEDs and measured redshifts the *actual* source structure played no role. Hence, we have in no way begged the question by having done so. Rather, this is precisely the chief advantage of simulating the data – it permits identifying the key signatures of the ICC halos and optimize the procedure to the detect them.

6. MONTE CARLO CONFIDENCE LEVEL ESTIMATES

The median angular power spectra presented in the previous section while indicative of the impact of different image features are poor representations of the angular power spectra associated with a single source realization. The fractional uncertainty in the power at a given multiple for a single image is order unity (see Appendix C.3), implying large deviations are typical. This is ameliorated by stacking the power spectra from multiple images as described in Section 5.5, which formally decreases the scatter by roughly $N_{\text{src}}^{-1/2}$. However, variations in the underlying intrinsic source and IGMF properties act to increase the fluctuations in \mathcal{P}_m .

Here we report the resulting 95% confidence-level regions for the various ICC halo models, and thereby IGMF models. The meaning of these is similar to that of a likelihood; they represent the probability of finding a given value of $\mathcal{P}_m^{\text{tot}}$ at a specified m for a given IGMF. As such, they present a natural way to assess the single realization afforded by the actual *Fermi* data.

Given the approximately bimodal nature of the ICC halos we have focused on two primary observables: the quadrupolar power $\mathcal{P}_2^{\text{tot}}$ and the sawtooth morphology. The latter is primarily in service to separating contributions to the angular power spectrum from ICC halos from other sources, including the potential sources of systematic error discussed in Section 5.4.2.

There are a number of additional potential observables. Typically, strong correlations between nearby multipoles limits the ability to leverage deviations for many m to improve statistical weight. Nevertheless, key systematics may be assessed using the distribution of $\mathcal{P}_m^{\text{tot}}$ at large m , set by the Poisson noise and

therefore the effective number of gamma rays used to construct the power spectrum estimate.

Thus, the essential experiment is to compare the low- m $\mathcal{P}_m^{\text{tot}}$ to the predictions of the various IGMF/halo models. Confidence in the simulation of the *Fermi* sources and their subsequent stacked angular power spectrum is obtained by comparing their large- m characteristics. The confidence with which any IGMF may be detected or excluded is then set by the degree to which the measured $\mathcal{P}_m^{\text{tot}}$ is inconsistent with the predicted value at, for example, $m = 2$.

The key theoretical input is then the anticipated *distribution* of $\mathcal{P}_m^{\text{tot}}$ for a given IGMF model. Here we describe how we estimate the ranges over which the angular power spectra can vary for a given halo model, and what these ranges are for the particular IGMF models described in Section 3.

6.1. Generating Mock Fermi Samples

The confidence with which a given IGMF can be detected or excluded based upon the gamma-ray image angular power spectra depends critically on a quantitative estimate of the range of, distribution of, and correlations among, the \mathcal{P}_m for the stacked power spectra of the *Fermi* sample. To assess this we perform a Monte Carlo sampling of the *Fermi* sample and its associated $\mathcal{P}_m^{\text{tot}}$. This synthesizes the steps described in Sections 3–5 for a given assumption regarding the geometry and strength of the IGMF.

In summary, for each realization of the *Fermi* sample we generate a set of realizations of each source in Table 1 consistent with their known intrinsic properties (e.g., N_{ph} and \mathcal{B}) and our estimates for the distributions of their unknown intrinsic properties (e.g., Γ_h , θ_{jet} , Θ , etc.). Within each image we choose the number of source and background photons from a Poisson distribution, taking N_{ph} and $\pi R_{\text{img}}^2 \mathcal{B}$ as means of the distribution, respectively (see Table 1). The source photons are distributed among the direct component and the ICC halo according to the particular halo model appropriate for the IGMF under consideration as described in Section 3.4. For large-scale, uniform IGMF models we further permit its local strength and orientation to vary as described in Section 4.4.

The resulting list of realized $\mathcal{P}_m^{\text{tot}}$ provide a proxy for the distribution of stacked angular power spectra associated with a given IGMF model. To quantify this distribution we quote one-sided 95% confidence-level regions, i.e., the values of $\mathcal{P}_m^{\text{tot}}$ that 95% of the realizations lie above, and independently the value that 95% lie below; because they are one-sided these are subtly different from the normal 2σ errors. This procedure of using one-sided confidence levels enables rejecting a model hypothesis at 95% confidence if it lies outside the one-sided half-interval. We also use the sample to explore the correlations between stacked angular power at different multipoles; when ICC halos contribute substantially to the $\mathcal{P}_m^{\text{tot}}$ the multipoles are generally strongly correlated, maintaining the sawtooth structure.

6.2. Anticipated Inverse Compton Halo Power Spectra

We now turn to the probability distributions of $\mathcal{P}_m^{\text{tot}}$ for various halo/IGMF models of interest. Note that in many cases the power at nearby multipoles are strongly correlated.

6.2.1. Null Case

We begin with the null hypothesis – the absence of any ICC halo component. In this case the image consists solely of the central point source and uniform background for each object in




Figure 16. Mock angular power spectra distributions when no ICC halos are present. The median value is shown by the filled pentagon, the triangles denote the one-sided 95% confidence interval. In comparison the black solid and dashed lines show the Poisson noise limit and its one-sided 95% confidence intervals, respectively, defined collectively by combining the distributions of all multipoles. This is produced with the 10^{-15} G, large-scale, uniform IGMF optimized source list.




Figure 17. Distribution of mock $\mathcal{P}_m^{\text{tot}}$ about the median for the first 4 multipoles ($m = 2, 3, 4, 5$ from red to blue) and the first 100 multipoles combined (black). For comparison a log-normal fit is shown (red dashed line) with mean -3.38 and standard deviation 0.129 . This is produced with the 10^{-15} G, large-scale, uniform IGMF optimized source list.

the *Fermi* sample. This represents the baseline against which all subsequent models that include ICC halos are compared. Here we use the 10^{-15} G, large-scale, uniform IGMF optimized source list.

The median angular power spectrum, including Front- and Back-converted events, and its 95% one-sided confidence limits are shown in Figure 16. As anticipated there is no structure in this case – the power spectrum is well fit by a constant corresponding to the effective Poisson noise limit. The distribution about the median values of $\mathcal{P}_m^{\text{tot}}$ are uncorrelated between multipoles and well fit by a log-normal distribution, shown in Figure 17.

We would expect any contaminating systematic (e.g., center-




Figure 18. Mock angular power spectra distributions when ICC halos associated with a small-scale, tangled IGMF are present. The median value is shown by the filled diamonds, the triangles denote the one-sided 95% confidence interval. In comparison the black solid and dashed lines show the Poisson noise limit and its one-sided 95% confidence intervals, respectively, defined collectively by combining the distributions of multipoles with $m \geq 50$. This is produced with the 10^{-15} G, large-scale, uniform IGMF optimized source list.




Figure 19. Pairwise correlations between the angular power in $m = 2, 3, 4$, and 5 multipoles for each mock realization of the *Fermi* sample when ICC halos associated with a small-scale, tangled IGMF. For reference the one-to-one line is shown by the red-dashed line. This is produced with the same set of realizations used to construct Figure 18, and utilizes the 10^{-15} G, large-scale, uniform IGMF optimized source list.

ing errors or dim neighbors) to contribute in a fashion similar to that seen for the single source power spectra presented in Section 5.4.2.

6.2.2. Small-scale, Tangled IGMF

The median and 95% confidence regions for the stacked angular power spectra for images containing ICC halos associ-

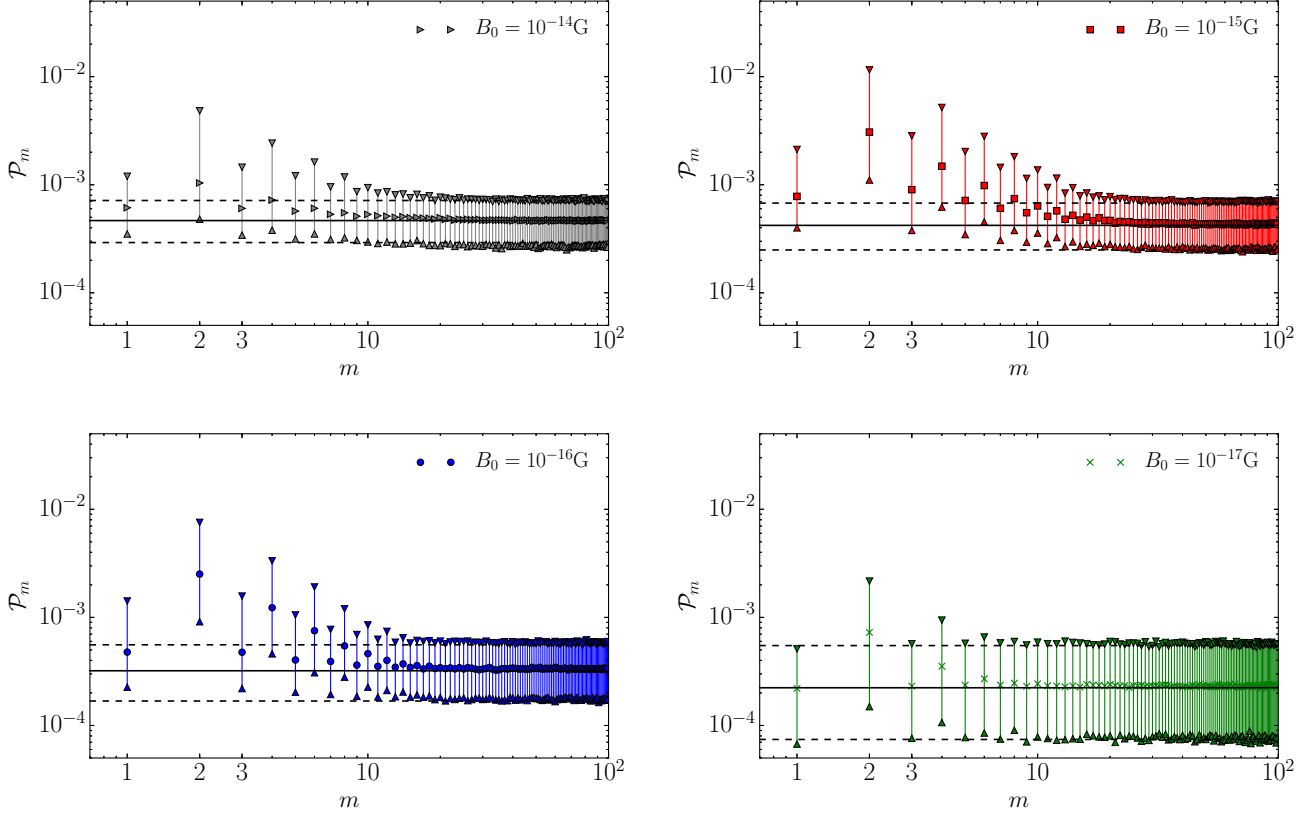


Figure 20. Mock angular power spectra distributions when ICC halos associated with a large-scale, uniform IGMF are present for $B_0 = 10^{-15}$ G (grey triangle, upper left), $B_0 = 10^{-15}$ G (red squares, upper right), 10^{-16} G (blue circles, lower left), and 10^{-17} G (green crosses, lower right). The median value is shown by the filled points, the triangles denote the one-sided 95% confidence interval. In comparison the black solid and dashed lines show the Poisson noise limit and its one-sided 95% confidence intervals, respectively, defined collectively by combining the distributions of multipoles with $m \geq 50$. Each case was produced using the associated optimized source list (see Table 1).

ated with a small-scale, tangled IGMF are shown in Figure 18. As anticipated by Figure 14 in comparison to Figure 16, and less directly by Figure 6, there is no discernable contribution from an ICC halo. This is a direct result of the combination of a low number of halo photons and the extreme foreshortening that results from the small viewing angles typical of the *Fermi* sources.

The cross-correlation of different multipoles, shown in Figure 19 reveals that for a small fraction of realizations ($< 5\%$) there are large deviations in the quadrupolar power. This suggests that alternative samples, composed of sources viewed at more oblique angles, may provide more promise for assessing the existence of ICC halos in the presence of a small-scale, tangled IGMF. We leave such an analysis for future work.

6.2.3. Large-scale, Uniform IGMF

Unlike the other models considered, the large-scale, uniform IGMF model has a free parameter in the form of the current IGMF strength. Thus, we show the distributions of stacked angular power for a variety of magnetic field strengths in Figure 20.

These all exhibit excess power at low- m , becoming less significant as the field strength nears 10^{-17} G, as anticipated from the single-image power spectra in Section 5. For $B_0 > 10^{-17}$ G the deviations from the Poisson noise are substantial, rising with increasing field strength; for $B_0 = 10^{-15}$ G the excess power extends to $m \approx 30$. Importantly, note that for

$B_0 = 10^{-15}$ G and 10^{-16} G the quadrupolar power is *inconsistent* with that for the null hypothesis, implying that the angular power spectrum will immediately distinguish between these two IGMF/halo models.

The breaking of the bipolar symmetry is clearly apparent in the stacked angular power, corresponding to the large power at odd multipoles. This is a natural consequence of the larger probability of viewing angles near the maximum permitted (see Section 3). Nevertheless, the sawtooth structure is directly evident in the medians of the $\mathcal{P}_m^{\text{tot}}$ for all cases.

Less obvious is that this remains true in individual realizations despite the large apparent variations implied by the extent of the 95% confidence regions; as seen in Figure 21, the power in all multipoles are strongly correlated, moving the sawtooth pattern up and down together. This is particularly true for the quadrupole, which overwhelmingly dominates all of the odd multipoles in more than 99% of the realizations, failing to do so only in the most pessimistic cases. This remains true for the Front- and Back-converted events independently, as shown in Figure 22, which shows the cross-correlation of the two event populations. *Thus, it is always possible to distinguish the large-scale, uniform IGMF model from the many potential sources of contamination in the angular power spectrum that show a smoothly falling excess of power at low- m generally.* The strong correlation among multipoles also means, however, that it remains possible that the power excess is low at all m , providing a fundamental limit on the minimum B_0 that can be




Figure 21. Pairwise correlations between the angular power in $m = 2, 3, 4$, and 5 multipoles for each mock realization of the *Fermi* sample when ICC halos associated with a large-scale, uniform IGMF with $B_0 = 10^{-15}$ G. For reference the one-to-one line is shown by the red-dashed line. This is produced with the same set of realizations used to construct the associated panel of Figure 20.

reliably differentiated from the null case.

6.3. Sensitivity to Source-Halo Confusion

How the potential confusion between source, halo and background photons is dealt with modifies the number of halo photons at large angular separations, and thus potentially the significance with which ICC halo contribution to the power spectrum may be identified. We therefore consider modifications to the scheme described in Section 3.4, extending the “source” region, i.e., the region within which photons will be assumed to contribute to the number of source photons to the 95% containment radius of the Pass 8R2_V6 ULTRACLEANVETO PSF. As discussed in Section 3.4 this extends to large angular radii, typically well beyond the point at which an extended component would be clearly visible. Nevertheless, as shown in Figure 23, this has little impact on the ability to identify ICC halos in the angular power spectrum.

6.4. Future Prospects with Fermi

The significance estimates shown in Figures 16, 18, and 20 are all for the current *Fermi* data set. However, the operational lifetime of *Fermi* may be considerably longer. Therefore, we also consider a gamma-ray data set that is doubled in size, i.e., the results of 16 years of *Fermi* observations. To facilitate a direct comparison and in the interest of simplicity we keep the source list identical.

As shown in Figure 24, the significance with which a large-scale, uniform IGMF can be detected or excluded is significantly improved. Of 10^3 realizations that assume an IGMF strength of $B_0 = 10^{-15}$ G, none were found that had a P_2 below the 95% confidence upper limit of the null model. Thus,




Figure 22. Pairwise cross-correlations between the angular power in $m = 2, 3, 4$, and 5 multipoles obtained from Front- (m_f) and Back-converted (m_b) events for each mock realization of the *Fermi* sample when ICC halos associated with a large-scale, uniform IGMF with $B_0 = 10^{-15}$ G. For reference the one-to-one line is shown by the red-dashed line. This is produced with the same set of realizations used to construct the associated panel of Figure 20.




Figure 23. Mock angular power spectrum assuming a large-scale, uniform IGMF with $B_0 = 10^{-15}$ G when the “source” region is 68% and 95% of the Pass 8R2_V6 ULTRACLEANVETO PSF, respectively. The median value is shown by the filled points, the triangles denote the one-sided 95% confidence intervals. In comparison the black solid and dashed lines show the Poisson noise limit and its one-sided 95% confidence intervals, respectively, defined collectively by combining the distributions of multipoles with $m \geq 50$. This is produced with the 10^{-15} G, large-scale, uniform IGMF optimized source list.

it would be possible to detect or exclude ICC halos associated with moderate IGMF strengths with significances well above 3σ .

More importantly, the multipoles for which this may be done extends beyond the quadrupole and includes the $m = 4$ mode. At $B_0 = 10^{-15}$ G this may be similarly excluded or detected with

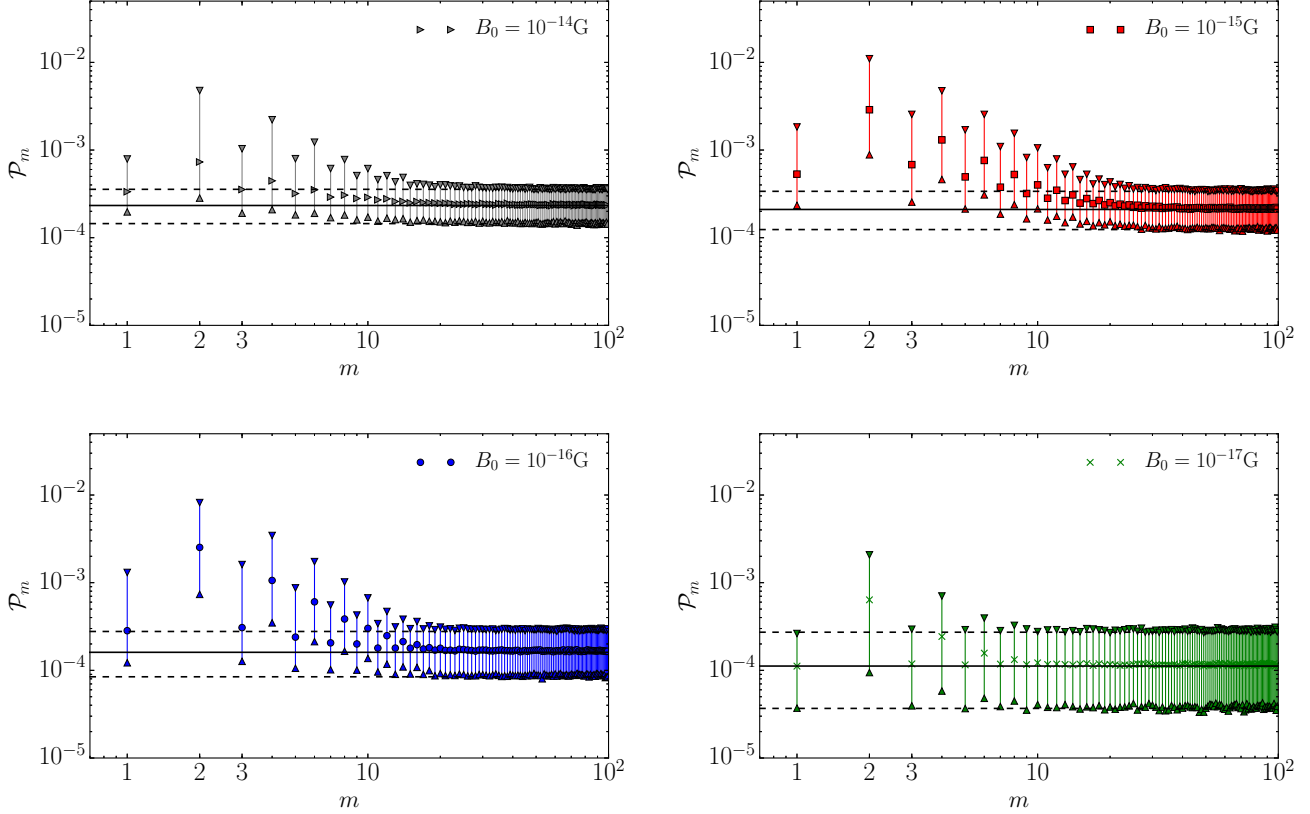


Figure 24. Mock angular power spectra distributions with double the number of photons (i.e., a projected total of 16 years of *Fermi* observations), when ICC halos associated with a large-scale, uniform IGMF are present for $B_0 = 10^{-14}$ G (grey triangle, upper left), $B_0 = 10^{-15}$ G (red squares, upper right), 10^{-16} G (blue circles, lower left), and 10^{-17} G (green crosses, lower right). The median value is shown by the filled points, the triangles denote the one-sided 95% confidence interval. In comparison the black solid and dashed lines show the Poisson noise limit and its one-sided 95% confidence intervals, respectively, defined collectively by combining the distributions of multipoles with $m \geq 50$. Each case was produced using the associated optimized source list (see Table 1).

significance $\gtrsim 3\sigma$; of 10^3 realizations 4 lay below the 95% confidence upper limit of the null model. As a result, should any of the systematic contaminants described in Section 5.4.2 be present (e.g., weak neighbors, unresolved backgrounds, etc.) the additional statistical power afforded by the increased fluence may prove critical to verifying any putative halo signal. However, the range of IGMF strengths that may be detected or excluded is not improved substantially.

For the null and small-scale, tangled IGMF halo models the result is primarily a reduced range of variation about a reduced Poisson noise limit; the ability to detect the latter is not substantially improved.

Note that this provides only a very pessimistic view of the value of an extended *Fermi* mission, even within the narrow confines of the detection of anisotropic ICC halos. A doubling of gamma-ray fluence not only increases the number of photons available in the sources listed in Table 1, but improves the significance with which structure can be identified in sources not listed in Table 1. As a result, the collection of *Fermi* sources that contribute positively to P_{det} will grow beyond those in Table 1 at all values of B_0 . As a result, the ability to detect ICC halos will benefit from an increased collection of sources, which we have not addressed, as well as the improved fluence.

7. CONCLUSIONS

In Broderick et al. (2016) we presented a semi-analytical computation of ICC halos and the generation of realizations around *Fermi* sources. There we found that the putative ICC halos that surround bright VHEGR sources are generally anisotropic. The origin of the anisotropy depends primarily on the structure of the IGMF in the AGNs vicinity, and otherwise weakly on its remaining intrinsic and observational properties. For small-scale, tangled IGMFs the halo structure is driven by the beamed nature of the VHEGR emission; for large-scale, uniform IGMFs the halo structure is a result of the opposite gyration of electrons and positrons coupled with the geometry of inverse Compton scattering in the high-energy limit.

In all cases, however, the structure of the gamma-ray map is bimodal. Here we suggest an explicit experiment to detect ICC halos using this bimodal structure that is insensitive to the many systematic uncertainties in the LAT instrumental response. This exploits the angular structure imposed by the bimodal structure instead of the radial structure often used. We propose to do this by generating one-dimensional angular power spectra, P_m , about each VHEGR source, following the application of an energy-dependent mask to remove the contribution of the direct emission from the central source. These have the considerable advantage that, unlike the images, they retain evidence of angular structure in the image when stacked regardless of absolute orientation.

The presence of bimodal structures in the underlying gamma-ray count maps generally produces a quadrupolar

power excess (i.e., at $m = 2$) and oscillating thereafter. The magnitude of the low-power excess and the depth of the oscillations depend on the intrinsic source parameters and the structure of the IGMF. Nevertheless, the sawtooth nature is generally present, both among different IGMF models and intrinsic parameter distributions as well as individual realizations of the images. Importantly, this is qualitatively distinct from the smooth power excesses induced by many potential observational and astronomical contaminants, including undetected neighbors, an unresolved point-source background, large background gradients, failures to center sources properly, and structure in the LAT instrument response.

Via Monte Carlo simulation of the images from VHEGR *Fermi* sources we have produced an optimized experiment for detecting ICC halos. That is, the ability to generate a large ensemble of realizations of images for the full sample enables the *theoretical* prediction of the anticipated *distribution* of the power at each m . We have therefore optimized both the procedure and sample of *Fermi* AGN. While the procedure is independent of the details underlying the ICC halo generation, the optimal list of *Fermi* sources is not; we report the optimized source samples for detecting ICC halos associated with large-scale, uniform IGMFs in Table 1. It is important to note that apart from the source SEDs and measured redshifts, this was performed *entirely* using *simulated* images, and thus is not colored by actual structures within the gamma-ray images.

Based upon this optimized experiment we are able to define confidence levels at which a given halo model can be detected or excluded using current and future *Fermi* data. We find that this depends on the origin of the ICC halo structure. ICC halos generated in the presence of a small-scale, tangled IGMF are unlikely to be detected using the *Fermi* AGN sample. This is a direct consequence of the extreme foreshortening associated

with the small viewing angles implied for most gamma-ray bright AGN and the intrinsic reduction in the ICC halo surface brightness associated with the isotropization of the generating pair population.

However, the ICC halos generated in the presence of a large-scale, uniform IGMF are detectable at high confidence for IGMF strengths between 10^{-17} G and 10^{-14} G, well matched to the range implied by gamma-ray observations of bright, nearby *Fermi* AGN. At weaker IGMF strengths the putative halos are sufficiently spatially concentrated that they are confused with the central source. At stronger IGMF strengths they are sufficiently spatially diffused that they are confused with neighboring sources. Were *Fermi* to continue operations for roughly eight more years the significance with which ICC halos could be detected or excluded would increase.

In a forthcoming companion publication we apply this to the existing sample of suitable *Fermi* blazars, placing constraints on the geometry and strength of the IGMF.

A.E.B. and M.S. receive financial support from the Perimeter Institute for Theoretical Physics and the Natural Sciences and Engineering Research Council of Canada through a Discovery Grant. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. C.P. acknowledges support by the European Research Council under ERC-CoG grant CRAGSMAN-646955 and by the Klaus Tschira Foundation. P.C. gratefully acknowledges support from the NASA ATP program through NASA grant NNX13AH43G, and NSF grant AST-1255469. A.L. receives financial support from an Alfred P. Sloan Research Fellowship, NASA ATP Grant NNX14AH35G, and NSF Collaborative Research Grant 411920 and CAREER grant 1455342. E.P. acknowledges support by the Kavli Foundation.

APPENDIX

A. GENERALIZED LAMBDA DISTRIBUTION FIT OF INTRINSIC OPENING ANGLES

The generalized lambda distribution (G λ D) provides a convenient parameterization of probability distributions (Ramberg & Schmeiser 1974). A key feature of the G λ D is that it parameterizes the *inverse* of the *cumulative* probability distribution instead of the probability density directly. This is particularly useful for generating random variates from the parameterized probability distribution.

The G λ D, with the Ramberg and Schmeiser (RS) parameterization, is defined by four constants, $\lambda_{1,2,3,4}$, in terms of which

$$F^{-1}(u) = \lambda_1 + \frac{u^{\lambda_3} - (1-u)^{\lambda_4}}{\lambda_2}, \quad (\text{A1})$$

where $u \in [0, 1]$ is the cumulative probability and $F^{-1}(u)$ is a map to the dependent variable. Drawing a random variate from the desired distribution is then reduced to choosing a uniform variate on $[0, 1]$ and evaluating $F^{-1}(u)$.

We performed a non-parametric, fit of the cumulative distribution of intrinsic radio opening angles, α_{int} from (Pushkarev et al. 2009), using the numerical statistics environment R (R Core Team 2014) and the G λ D RS fitting package by GLDEX (Su 2016, 2007) and performed a maximum likelihood fit. A KS test was performed to analyze the quality of the fit and the result was a p-value 0.9704 i.e. consistent with the assumption that they are drawn from the same distribution. The four G λ D parameters are $\lambda_1 = 0.8261001$, $\lambda_2 = -0.63510851$, $\lambda_3 = -0.07892053$, $\lambda_4 = -0.2946635$, and the resulting probability distribution is shown in Figure 9.

B. MAXIMUM-LIKELIHOOD CENTER FINDING

Here we describe how we produce refined high-accuracy estimates of the source location. We begin with a parameterized model for the probability distribution of photons in an image that describes a Gaussian source on a uniform background. This has three parameters: the source position, μ , the ratio of photons in the background to those in the source, A , and the width of the source distribution, σ . We will assume that the last is fixed by the characteristic size of the *Fermi* PSF. Finally, and critically, we will only perform the center-fitting within a known circular window with angular radius R about some fixed initial position. Thus, the

probability of a given photon being located at position \mathbf{x} given parameter $\mathbf{p} = (\boldsymbol{\mu}, A)$ is

$$P(\mathbf{x}|\mathbf{p}) = \Theta(R - |\mathbf{x}|)N \left[\frac{1}{2\pi\sigma^2} e^{-(\mathbf{x}-\boldsymbol{\mu})^2/2\sigma^2} + \frac{A}{\pi R^2} \right], \quad (\text{B1})$$

These are not all independent, related by the normalization condition upon $P(\mathbf{x}|\mathbf{p})$:

$$1 = \int d^2x P(\mathbf{x}|\mathbf{p}) = \int_0^R r dr \int_0^{2\pi} d\theta N \left[\frac{1}{2\pi\sigma^2} e^{-(r^2 - 2r\mu \cos\theta + \mu^2)/2\sigma^2} + \frac{A}{\pi R^2} \right] = NA + Ne^{-\mu^2/2\sigma^2} \int_0^{R/\sigma} dx x e^{-x^2/2} I_0(xm), \quad (\text{B2})$$

where $x \equiv r/\sigma$, $m \equiv \mu/\sigma$, and $I_0(x)$ is the modified Bessel function of the first kind. This final integral is difficult to perform in practice due to the presence of I_0 . It maybe be simplified via various asymptotic expansions, though in practice, it is most easily evaluated numerically.

However it is computed, the resulting integral is a function of R^2 , σ^2 , and μ^2 only, and thus we define:

$$I(R^2, \sigma^2, \mu^2) \equiv e^{-\mu^2/2\sigma^2} \int_0^{R/\sigma} dx x e^{-x^2/2} I_0(xm). \quad (\text{B3})$$

Regardless how $I(R^2, \sigma^2, \mu^2)$ is computed, the normalization condition becomes

$$N = \frac{1}{A + I(R^2, \sigma^2, \mu^2)}, \quad (\text{B4})$$

and thus, our probability distribution is

$$P(\mathbf{x}|\mathbf{p}) = \Theta(R - |\mathbf{x}|) \frac{(\pi R^2/2\pi\sigma^2) e^{-(\mathbf{x}-\boldsymbol{\mu})^2/2\sigma^2} + A}{\pi R^2 [A + I(R^2, \sigma^2, \mu^2)]}. \quad (\text{B5})$$

The corresponding likelihood for a set of $\{\mathbf{x}_j\}$ is then

$$L(\{\mathbf{x}_j\}|\mathbf{p}) = -2 \sum_j \log \left\{ \frac{(\pi R^2/2\pi\sigma^2) e^{-(\mathbf{x}_j-\boldsymbol{\mu})^2/2\sigma^2} + A}{\pi R^2 [A + I(R^2, \sigma^2, \mu^2)]} \right\} = -2 \sum_j \log \left(\frac{\Delta_j + A}{I + A} \right) + \text{const.}, \quad (\text{B6})$$

where we ignore the constant πR^2 term and define

$$\Delta_j = \frac{\pi R^2}{2\pi\sigma^2} e^{-(\mathbf{x}_j-\boldsymbol{\mu})^2/2\sigma^2}. \quad (\text{B7})$$

Finally, to avoid issues with A dropping below zero, we replace it with e^α , which is positive definite. Thus,

$$L = -2 \sum_j \log \left(\frac{\Delta_j + e^\alpha}{I + e^\alpha} \right) + \text{const.} \quad (\text{B8})$$

We supplement these with the derivatives:

$$\frac{\partial L}{\partial \mu} = -2 \sum_j \left[\frac{\Delta_j}{\Delta_j + e^\alpha} \frac{(\mathbf{x}_j - \boldsymbol{\mu})}{\sigma^2} - \frac{2\boldsymbol{\mu}}{I + e^\alpha} \frac{\partial I}{\partial \mu^2} \right] \quad \text{and} \quad \frac{\partial L}{\partial \alpha} = -2 \sum_j e^\alpha \left[\frac{1}{\Delta_j + e^\alpha} - \frac{1}{I + e^\alpha} \right]. \quad (\text{B9})$$

The values of $\boldsymbol{\mu}$ and α are then obtained via numerically minimizing L .

C. EXAMPLE ANGULAR POWER SPECTRA

Here we present a number of angular power spectra for toy models. These are primarily for illustrative and contextual purposes – power spectra in the text are obtained for physically motivated mock images in various circumstances. In all cases we make use of the following general definitions. The angular power associated with a given image is

$$\mathcal{P}_m = \frac{1}{N^2} \left| \sum_j e^{im\theta_j} \right|^2 = \frac{1}{N^2} \sum_{j,k} e^{im(\theta_j - \theta_k)} = \frac{1}{N} + \frac{1}{N^2} \sum_{j \neq k} e^{im(\theta_j - \theta_k)}. \quad (\text{C1})$$

Note that generally, for $m=0$ we have $\mathcal{P}_0 = N^{-2} \sum_{j,k} 1 = 1$ independent of the details of the distribution of θ_j .

To compute the statistical properties of \mathcal{P}_m for $m > 0$ we must specify the probability distribution of the $\theta_{N,m}$, which we describe via the probability density of finding a photon at θ , $\wp(\theta)$. (The independence of $m=0$ is a consequence of the normalization condition on $\wp(\theta)$.) Hence, assuming each photon is independent, the mean is given by

$$\langle \mathcal{P}_m \rangle = \int d\theta_1 \dots d\theta_N \wp(\theta_1) \dots \wp(\theta_N) \left[\frac{1}{N} + \frac{1}{N^2} \sum_{j \neq k} e^{im(\theta_j - \theta_k)} \right] = \frac{1}{N} + \frac{N-1}{N} \int d\theta_1 d\theta_2 \wp(\theta_1) \wp(\theta_2) e^{im(\theta_1 - \theta_2)} = \frac{1}{N} + \frac{N-1}{N} |\varphi_m|^2 \quad (\text{C2})$$

where φ_m is the characteristic function of the photon distribution:

$$\varphi_m \equiv \int d\theta \wp(\theta) e^{im\theta}. \quad (\text{C3})$$

The variance may be computed in a similar way, though some care must be taken with counting the degeneracies of the various terms, ultimately yielding for $m > 0$

$$\begin{aligned} \sigma_{\mathcal{P}_m}^2 &= \langle \mathcal{P}_m^2 \rangle - \langle \mathcal{P}_m \rangle^2 \\ &= \frac{1}{N^4} \int d\theta_1 \dots d\theta_N \wp(\theta_1) \dots \wp(\theta_N) \left[N^2 + 2N \sum_{j \neq k} e^{im(\theta_j - \theta_k)} + \sum_{j \neq k, a \neq b} e^{im(\theta_j - \theta_k + \theta_a - \theta_b)} \right] - \left[N + N(N-1) |\phi_m|^2 \right]^2 \\ &= \frac{1}{N^4} \int d\theta_1 \dots d\theta_N \wp(\theta_1) \dots \wp(\theta_N) \left[N^2 + 2N \sum_{j \neq k} e^{im(\theta_j - \theta_k)} + \sum_{j=b \neq k=a} 1 + \sum_{j=b \neq k \neq a} e^{im(\theta_a - \theta_k)} + \sum_{j \neq b \neq k=a} e^{im(\theta_j - \theta_b)} \right. \\ &\quad \left. + \sum_{j=a \neq k=b} e^{im(2\theta_j - 2\theta_k)} + \sum_{j=a \neq k \neq b} e^{im(2\theta_j - \theta_k - \theta_b)} + \sum_{j \neq a \neq k=b} e^{im(\theta_j - 2\theta_k + \theta_a)} + \sum_{j \neq k \neq a \neq b} e^{im(\theta_j - \theta_k + \theta_a - \theta_b)} \right] - \left[N + N(N-1) |\phi_m|^2 \right]^2 \quad (\text{C4}) \\ &= \frac{1}{N^4} \left[N^2 + 2N^2(N-1) |\varphi_m|^2 + N(N-1) + 2N(N-1)(N-2) |\varphi_m|^2 + N(N-1) |\varphi_{2m}|^2 \right. \\ &\quad \left. + N(N-1)(N-2) \left(\varphi_{2m}^* \varphi_m^2 + \varphi_{2m} \varphi_m^{*2} \right) + N(N-1)(N-2)(N-3) |\varphi_m|^4 - N^2 - 2N^2(N-1) |\varphi_m|^2 - N^2(N-1)^2 |\varphi_m|^4 \right] \\ &= \frac{N-1}{N^3} \left[1 + (2N-4) |\varphi_m|^2 - (4N-6) |\varphi_m|^4 + |\varphi_{2m}|^2 + (N-2) \left(\varphi_{2m}^* \varphi_m^2 + \varphi_{2m} \varphi_m^{*2} \right) \right]. \end{aligned}$$

Identifying $\langle \mathcal{P}_m \rangle$ and $\sigma_{\mathcal{P}_m}$ for specific models now requires the specification of φ_m . We consider a handful of specific cases below.

C.1. Isotropic Background

We begin with an isotropic background, i.e.,

$$\wp(\theta) = \frac{1}{2\pi} \quad \Rightarrow \quad \varphi_m = \delta_{m0}, \quad (\text{C5})$$

and thus for $m > 0$

$$\langle \mathcal{P}_m \rangle = \frac{1}{N} \quad \text{and} \quad \sigma_{\mathcal{P}_m}^2 = \frac{N-1}{N^3}, \quad (\text{C6})$$

Note that this does not imply a uniform background, only one that is isotropic about the chosen image center. This is simply the Poisson noise limit, and as is typical in power spectrum estimation, based on a single source the variance in the estimate of \mathcal{P}_m is comparable to its mean.

C.2. Offset Gaussian

We now consider a peaked source offset from the origin, as may occur if the source location estimate is in error. Here we make simplifying assumption that the source Gaussian, i.e., set the 2-dimensional photon probability density to be

$$\wp(\alpha, \delta) = \frac{1}{2\pi\sigma^2} e^{-[(\alpha - \eta\sigma)^2 + \delta^2]/2\sigma^2} = \frac{1}{2\pi\sigma^2} e^{-(r^2 + \eta^2\sigma^2)/2\sigma^2} \left(1 + \frac{\eta r \cos \theta}{\sigma} + \frac{\eta^2 r^2 \cos^2 \theta}{2\sigma^2} + \frac{\eta^3 r^3 \cos^3 \theta}{6\sigma^3} + \dots \right). \quad (\text{C7})$$

Integrating over radius, this gives the needed angular probability density:

$$\wp(\theta) = \int r dr \wp(\alpha, \delta) = \frac{1}{2\pi} \left[1 + \sqrt{\frac{\pi}{2}} \eta \cos \theta + \eta^2 \cos^2 \theta + \sqrt{\frac{\pi}{8}} \eta^3 \cos^3 \theta + \dots \right]. \quad (\text{C8})$$

Therefore, the characteristic function is

$$\varphi_m = \delta_{m0} + \frac{1}{2} \sqrt{\frac{\pi}{2}} \eta (\delta_{m1} + \delta_{m-1}) + \frac{1}{4} \eta^2 (\delta_{m2} + 2\delta_{m0} + 2\delta_{m-2}) + \frac{1}{8} \sqrt{\frac{\pi}{8}} \eta^3 (\delta_{m3} + 3\delta_{m1} + 3\delta_{m-1} + \delta_{m-3}) + \dots \quad (\text{C9})$$

To lowest order in η at each m , this gives

$$\langle \mathcal{P}_m \rangle = N + N(N-1) \left[\delta_{m0} + \frac{\pi}{8} \eta^2 \delta_{m1} + \frac{1}{16} \eta^4 \delta_{m2} + \frac{\pi}{512} \eta^6 \delta_{m3} + \dots \right]. \quad (\text{C10})$$

Note that the contribution to the high- m multipoles falls off as η^{2m} . Not surprisingly, the variance also falls off quite rapidly in η :

$$\sigma_{\mathcal{P}_m}^2 \simeq \sum_j 2N^3 \left[\delta_{m0} + \frac{\pi}{8} \eta^2 \delta_{m1} + \frac{1}{16} \eta^4 \delta_{m2} + \frac{\pi}{512} \eta^6 \delta_{m3} + \dots \right]. \quad (\text{C11})$$

We can construct a general formula, using the integral

$$\int_0^\infty dr r^{n+1} e^{-r^2/2} = 2^{n/2} \Gamma\left(\frac{n}{2} + 1\right), \quad (\text{C12})$$

with which we have

$$\mathcal{P}(\theta) = \frac{1}{2\pi} \sum_{j=0}^\infty \frac{2^{j/2}}{j!} \Gamma\left(\frac{j}{2} + 1\right) \eta^j \cos^j \theta \quad \Rightarrow \quad \varphi_m = \sum_{j=0}^\infty \frac{2^{-j/2}}{j!} \Gamma\left(\frac{j}{2} + 1\right) \eta^j \sum_{k=-j}^j \frac{j!}{k!(j-k)!} \delta_{mk}, \quad (\text{C13})$$

in terms of which we have the lowest-order terms in η for each m in the power spectrum

$$\langle \mathcal{P}_m \rangle \approx |\varphi_m|^2 \approx 2^{-m} \left[\frac{1}{m!} \Gamma\left(\frac{m}{2} + 1\right) \right]^2 \eta^{2m} < 2^{-3m/2} \eta^{2m} \quad \text{and} \quad \sigma_{\mathcal{P}_m}^2 \approx 2 |\phi_m|^2 < 2^{1-3m/2} \eta^{2m}. \quad (\text{C14})$$

C.3. Weak Image Components

There are many instances where the image is comprised of a dominant, isotropic component and sub-dominant structured component. In this case we generally have

$$\mathcal{P}(\theta) = \frac{1-\eta}{2\pi} + \eta h(\theta) \quad \Rightarrow \quad \varphi_m = \delta_{m0} + \eta \psi_m, \quad (\text{C15})$$

where ψ_m is the structure function associated with the unit-normalized $h(\theta)$ perturbing distribution, with strength described by order parameter $\eta \ll 1$ that is equal to the ratio of the photons associated with the perturbation to those in the isotropic background. Then, for $m > 0$

$$\langle \mathcal{P}_m \rangle \approx \frac{1}{N} + \eta^2 |\psi_m|^2 \quad \text{and} \quad \sigma_{\mathcal{P}_m}^2 \approx \frac{1}{N^2} + \frac{2}{N} \eta^2 |\psi_m|^2. \quad (\text{C16})$$

Note that again $\sigma_{\mathcal{P}_m} \approx \mathcal{P}_m$, which motivates combining multiple estimates of \mathcal{P}_m to obtain improved statistical significance of any detection.

C.3.1. Gaussian Neighbor

Distant peaked sources, arising from unidentified neighbors, will also induce dipolar power. Again we make the simplifying assumption that the neighboring source is Gaussian, which if it is sufficiently far is approximately described by

$$h(\theta) = \frac{e^{-(\theta-\theta_0)^2/2w^2}}{\sqrt{2\pi}w} \quad \Rightarrow \quad \psi_m = e^{im\theta_0 - w^2 m^2/2}, \quad (\text{C17})$$

where θ_0 and $w = \sigma/\Delta r$ are the location and width of the companion in polar angle measured about the central source. The associated mean power spectrum is

$$\langle \mathcal{P}_m \rangle \approx \frac{1}{N} + \eta^2 e^{-w^2 m^2} \quad \text{and} \quad \sigma_{\mathcal{P}_m}^2 \approx \frac{1}{N^2} + \frac{2}{N} \eta^2 e^{-w^2 m^2}. \quad (\text{C18})$$

As anticipated this generates substantial dipolar power, smoothly falling off on a scale inverse set by the width.

C.3.2. Bimodal Gaussian

As an approximation to what the angular power spectrum from a bi-modal halo may look like we consider a bimodal Gaussian, consisting of two sources located 180° apart. That is,

$$h(\theta) = \frac{1}{\sqrt{8\pi}w} \left(e^{-(\theta-\theta_0)^2/2w^2} + e^{-(\theta-\theta_0+\pi)^2/2w^2} \right) \quad \Rightarrow \quad \psi_m = \Delta_m e^{im\theta_0 - w^2 m^2/2}, \quad (\text{C19})$$

where Δ_m is unity for even m and zero for odd m . This is similar to the Gaussian neighbor, with the exception of the factor of Δ_m that arises from the beating between the two sources. Hence, the mean power spectrum is,

$$\langle \mathcal{P}_m \rangle \approx \frac{1}{N} + \eta^2 \Delta_m e^{-w^2 m^2} \quad \text{and} \quad \sigma_{\mathcal{P}_m}^2 \approx \frac{1}{N^2} + \frac{2}{N} \eta^2 \Delta_m e^{-w^2 m^2}. \quad (\text{C20})$$

Note that this exhibits the clear sawtooth indicative of the bimodal structures generally.

D. ESTIMATE OF PSF-INDUCED ANGULAR STRUCTURE

We make a quantitative estimate of the impact of the intrinsic anisotropy of the *Fermi* LAT PSF. To do this we follow the description in §5.2.3 of Ackermann et al. (2012a) and adopt an azimuthally dependence for the PSF of the form

$$f(\theta) = 1 + q_0 \xi^{q_1} \quad \text{where} \quad \xi = \frac{4}{\pi} \left| \left(\theta \bmod \frac{\pi}{2} \right) - \frac{\pi}{4} \right|. \quad (\text{D1})$$

where q_0 , q_1 , and ξ have the same meaning as in Equation (15) of Ackermann et al. (2012a). The impact of this is to modify the distribution of photons in the image, inducing a small $m = 4$ perturbation to the angular probability distribution.

The gamma rays in a typical image from a bright *Fermi* AGN are accumulated over many years, and therefore the LAT presents different orientations. Were the orientation completely independent for each photon, the net PSF would become isotropic and therefore any remaining structure would be absent. However, we may imagine that subsequent photons have LAT orientations that are closely correlated. The roll angle of the LAT changes primarily as a result of the fixed telescope orientation relative to the Sun. The square geometry of the LAT implies that the angular impact on PSF is effectively averaged by the time it has rotated by 45° , corresponding to roughly 45 days. Over this timescale a non-varying *Fermi* source would accumulate less than 2% of its total fluence or roughly 100 photons.

Thus, we estimate the residual power from $n \approx 50$ epochs of $N \approx 100$ events, assuming perfect correlation in the LAT orientation for the latter. Within each subgroup there is a fixed roll angle, ϑ_μ , and hence relative to the orientation of the LAT the polar angle of the photons are $\theta_{j,\mu} = \theta_j - \vartheta_\mu$. Between subgroups the roll angle is effectively uncorrelated. Thus, for $m > 0$,

$$\langle \mathcal{P}_m \rangle = \frac{1}{(nN)^2} \sum_{\mu,\nu} \sum_{j,k} \langle e^{im(\theta_j - \theta_k)} e^{-im(\vartheta_j + \vartheta_k)} \rangle = \frac{1}{nN^2} \sum_{j,k} \langle e^{im(\theta_j - \theta_k)} \rangle, \quad (\text{D2})$$

where the independence of different ϑ_μ was employed.

Following the prefactor of n^{-1} is simply the angular power spectrum for the single epoch, for which the photon distribution is perturbed from isotropy. This is precisely the case considered in Appendix C.3. Here, the perturbation is $h(\theta) = \xi/\pi$, where we have set $q_1 = 1$ for simplicity and implicitly subsumed q_0 into $\eta = 2q_0$. The associated characteristic function is

$$\psi_m = \int d\theta \frac{\xi}{\pi} e^{im\theta} = \left(1 + e^{im\pi/2} + e^{im\pi} + e^{3im\pi/2} \right) \int_0^{\pi/2} d\theta \frac{\xi}{\pi} e^{im\theta} = -\frac{i}{2} \Sigma_m e^{im\pi/4} j_1 \left(\frac{m\pi}{4} \right), \quad (\text{D3})$$

where Σ_m is unity for m that are divisible by 4 and zero otherwise, and $j_1(x)$ is the spherical Bessel function of the first kind of order 1. Therefore, the corresponding mean angular power spectrum is

$$\langle \mathcal{P}_m \rangle = \frac{1}{n} \left[\frac{1}{N} + \eta^2 \Sigma_m \frac{1}{4} j_1^2 \left(\frac{m\pi}{4} \right) \right] \lesssim \frac{1}{nN} + \frac{\eta^2}{n} \frac{\Sigma_m}{4\pi^2 m^2}. \quad (\text{D4})$$

The excess is non-vanishing only for $m = 4$ and its harmonics; in particular there is no power at $m = 2$. For a typical value of $\eta \approx 0.2$, corresponding to $q_0 \approx 0.1$, the excess at $m = 4$ is $\lesssim 10^{-6}$ which is much smaller than the $(nN)^{-1} \gtrsim 10^{-4}$.

In practice, the angular power associated with structure in the PSF is further reduced by the fraction of the total photons those from the source compose. This is impacted both by the surface brightness of the background, typically contributing 25% for the brightest sources, and the source mask, typically cutting the source contribution by half. This makes only an order unity correction.

E. SOURCE SAMPLE

Here we collect the images and fits for the 18 sources listed in Table 1 that comprise the full set of sources within the optimized source lists. Details of the fitting process are described in Section 4.1; we only summarize these here.

There were 84 sources selected on the basis of their 2FHL 50 GeV–2 TeV flux, as well as having a redshift below 0.5 and a 3LAC counterpart. Gamma-ray SEDs were produced by compiling the 0.1–0.3 GeV, 0.3–1 GeV, 1–3 GeV, 3–10 GeV, and 10–100 GeV flux measurements reported in the 3LAC and 50–171 GeV, 171–585 GeV, and 585–2000 GeV reported in the 2FHL. These were deabsorbed via Equation 11 evaluated at the geometric center of the energy bin. A maximum-likelihood fit of the broken-power-law SED model was performed to each candidate source, described by low- and high-energy spectral indexes, Γ_l , Γ_h , a pivot energy E_p , and a normalization.

Sources for which the source-frame TeV flux, i.e., $E^2 dN/dE$ at $(1+z)$ TeV, was above 10^{-9} GeV cm $^{-2}$ s $^{-1}$, were then visually inspected for neighboring sources within 2° or large background gradients. Sufficiently isolated, TeV-bright sources comprise the source sample employed here, consisting of the 27 objects listed in Tables 2 and 3. Note that the sources that comprise the optimized source lists are chosen from these 27 after simulating the ability to detect ICC halos (see Section 5.6).

The observed and deabsorbed SEDs and our fits for the 18 sources used here are shown in Figures 26 and 27.

REFERENCES

- | | |
|--|---|
| <p>Abdo, A. A., et al. 2010a, ApJS, 188, 405
 —. 2010b, PhRvL, 104, 101101
 —. 2010c, ApJ, 716, 30
 Acero, F., et al. 2015, ApJS, 218, 23
 Ackermann, M., et al. 2011, ApJ, 743, 171
 —. 2012a, ApJS, 203, 4
 —. 2012b, Sci, 338, 1190</p> | <p>—, 2013a, ApJ, 765, 54
 —. 2013b, ArXiv:1306.6772
 —. 2015a, ApJ, 799, 86
 —. 2015b, ApJ, 810, 14
 —. 2016a, ApJS, 222, 5
 —. 2016b, PhRvL, 116, 151105
 Albert, J., et al. 2007, ApJ, 663, 125
 Alexander, P., & Leahy, J. P. 1987, MNRAS, 225, 1
 Ando, S., & Kusenko, A. 2010, ApJL, 722, L39</p> |
|--|---|




Figure 25. The rank-ordered list of TeV fluxes. The nominal cutoff is shown by the dotted line. Sources are ranked by the 1σ lower limit on the TeV fluxes. Sources are color coded by SED type: convex ($\Gamma_h > \Gamma_l$) and concave ($\Gamma_h < \Gamma_l$) SEDs are shown in red and blue, respectively.

Table 2
List of Sources

3FGL	2FHL	Names		Class ^a	z^b	Fluxes	
		TeV	Other			$F_{3\text{FGL}}^c$ ($\text{ph cm}^{-2}\text{s}^{-1}$)	$F_{2\text{FHL}}^d$ ($\text{ph cm}^{-2}\text{s}^{-1}$)
J1104.4+3812	J1104.4+3812	J1104+3811	Mkn 421	bll	0.031	3.03×10^{-8}	1.24×10^{-9}
J1653.9+3945	J1653.9+3945	J1653+3945	Mkn 501	bll	0.0337	9.74×10^{-9}	4.78×10^{-10}
J0449.4-4350	J0449.4-4349	J0449-4350	PKS 0447-439	bll	0.205	1.03×10^{-8}	1.65×10^{-10}
J0508.0+6736	J0507.9+6737	J0507+6737	1ES 0502+675	bll	0.34	2.12×10^{-9}	1.63×10^{-10}
J1015.0+4925	J1015.0+4926	J1015+4926	1ES 1011+496	bll	0.212	7.35×10^{-9}	1.62×10^{-10}
J2000.0+6509	J2000.1+6508	J1959+6508	1ES 1959+650	bll	0.047	5.84×10^{-9}	1.57×10^{-10}
J0650.7+2503	J0650.7+2502	J0650+2503	1ES 0647+250	bll	0.203	2.16×10^{-9}	1.26×10^{-10}
J2009.3-4849	J2009.4-4849	J2009-4849	PKS 2005-489	bll	0.071	3.55×10^{-9}	9.70×10^{-11}
J0303.4-2407	J0303.3-2407	J0303-2407	PKS 0301-243	bll	0.26	5.74×10^{-9}	8.62×10^{-11}
J1444.0-3907	J1443.9-3909	—	PKS 1440-389	bll	0.0654	3.23×10^{-9}	8.22×10^{-11}
J2347.0+5142	J2347.1+5142	J2346+5142	1ES 2344+514	bll	0.044	2.42×10^{-9}	7.48×10^{-11}
J0648.8+1516	J0648.6+1516	J0648+1516	RX J0648.7+1516	bll	0.179	1.45×10^{-9}	5.62×10^{-11}
J0543.9-5531	J0543.9-5533	—	1RXS J054357.3-55320	bll	0.273	1.61×10^{-9}	5.27×10^{-11}
J1512.8-0906	J1512.7-0906	J1512-0906	PKS 1510-08	fsrq	0.36	4.11×10^{-8}	4.59×10^{-11}
J1120.8+4212	J1120.8+4212	—	RBS 0970	bll	0.124	1.16×10^{-9}	4.18×10^{-11}
J1117.0+2014	J1116.9+2014	—	RBS 0958	bll	0.138	1.74×10^{-9}	3.85×10^{-11}
J2250.1+3825	J2249.9+3826	J2250+3824	B3 2247+381	bll	0.119	1.10×10^{-9}	3.38×10^{-11}
J0627.0-3529	J0626.9-3528	—	PKS 0625-35	rdg	0.05494	1.43×10^{-9}	2.92×10^{-11}
J0622.4-2606	J0622.4-2604	—	PMN J0622-2605	bll	0.41449	1.14×10^{-9}	2.58×10^{-11}
J2131.5-0915	J2131.4-0914	—	RBS 1752	bll	0.449	8.63×10^{-10}	2.49×10^{-11}
J2016.4-0905	J2016.5-0904	—	PMN J2016-0903	bll	0.367	1.40×10^{-9}	1.86×10^{-11}
J0115.8+2519	J0115.8+2519	—	RX J0115.7+2519	bll	0.358	1.15×10^{-9}	1.83×10^{-11}
J0757.0+0956	J0756.8+0955	—	PKS 0754+100	bll	0.266	1.86×10^{-9}	1.73×10^{-11}
J0958.6+6534	J0958.3+6535	—	S4 0954+65	bll	0.367	1.38×10^{-9}	1.54×10^{-11}
J2329.2+3754	J2329.2+3754	—	NVSS J232914+375414	bll	0.264	8.40×10^{-10}	1.43×10^{-11}
J1436.8+5639	J1437.0+5639	—	RBS 1409	bll	0.15	4.71×10^{-10}	1.40×10^{-11}
J1442.8+1200	J1442.9+1159	J1442+1200	1ES 1440+122	bll	0.16309	5.59×10^{-10}	1.39×10^{-11}

^a Source class type: bll, fsrq, and rdg correspond to BL Lac, flat-spectrum radio quasars, and radio galaxy, respectively

^b 3FGL value

^c 1-100 GeV fluence from 3FGL

^d 50-2000 GeV fluence from 3FGL

Table 3
Source SED Fit Parameters and TeV Flux

3FGL Name	F_{TeV}^a (GeV cm ⁻² s ⁻¹)	f_0 (ph cm ⁻² s ⁻¹)	E_p (GeV)	Γ_l	Γ_h^b
Convex Spectrum Sources ($\Gamma_h > \Gamma_l$)					
J1104.4+3812	$4.69^{+1.20}_{-1.04} \times 10^{-8}$	6.53×10^{-8}	95.38	1.77	$2.14^{+0.11}_{-0.10}$
J1653.9+3945	$1.14^{+0.79}_{-0.58} \times 10^{-8}$	3.21×10^{-8}	236.30	1.72	$2.72^{+0.51}_{-0.37}$
J0449.4-4350	$5.07^{+2.14}_{-1.75} \times 10^{-9}$	1.58×10^{-8}	23.03	1.81	$2.30^{+0.11}_{-0.10}$
J1015.0+4925	$1.19^{+0.23}_{-0.21} \times 10^{-8}$	7.89×10^{-9}	2.71	1.75	$1.93^{+0.03}_{-0.03}$
J2000.0+6509	$4.49^{+4.91}_{-3.45} \times 10^{-9}$	1.15×10^{-8}	658.00	1.87	$4.25^{+1.58}_{-1.60}$
J0650.7+2503	$6.10^{+20.21}_{-5.23} \times 10^{-9}$	8.75×10^{-9}	281.60	1.67	$2.29^{+1.64}_{-1.35}$
J2009.3-4849	$1.15^{+1.89}_{-0.92} \times 10^{-9}$	7.08×10^{-9}	76.45	1.77	$2.71^{+0.64}_{-0.39}$
J0303.4-2407	$7.03^{+1.44}_{-1.28} \times 10^{-9}$	5.68×10^{-9}	1.02	1.78	$1.97^{+0.03}_{-0.03}$
J1444.0-3907	$2.94^{+1.44}_{-1.18} \times 10^{-9}$	5.43×10^{-9}	11.90	1.67	$2.14^{+0.13}_{-0.09}$
J2347.0+5142	$4.47^{+1.39}_{-1.21} \times 10^{-9}$	2.43×10^{-9}	1.81	1.69	$1.90^{+0.05}_{-0.04}$
J0648.8+1516	$4.49^{+1.74}_{-1.38} \times 10^{-9}$	1.21×10^{-9}	1.03	0.97	$1.81^{+0.05}_{-0.05}$
J0543.9-5531	$5.90^{+1.88}_{-1.64} \times 10^{-9}$	1.17×10^{-9}	0.84	0.46	$1.77^{+0.05}_{-0.04}$
J1512.8-0906	$1.71^{+0.27}_{-0.21} \times 10^{-9}$	7.24×10^{-8}	1.27	2.22	$2.56^{+0.02}_{-0.02}$
J1120.8+4212	$1.70^{+1.61}_{-1.07} \times 10^{-9}$	2.92×10^{-9}	28.65	1.56	$2.15^{+0.29}_{-0.19}$
J1117.0+2014	$2.53^{+1.33}_{-1.03} \times 10^{-9}$	1.89×10^{-9}	3.02	1.77	$1.95^{+0.09}_{-0.08}$
J2016.4-0905	$1.63^{+1.00}_{-0.74} \times 10^{-9}$	1.50×10^{-9}	2.37	1.89	$1.99^{+0.11}_{-0.11}$
Concave Spectrum Sources ($\Gamma_h < \Gamma_l$)					
J0508.0+6736	$4.09^{+0.74}_{-0.70} \times 10^{-8}$	1.32×10^{-9}	2.32	1.81	$1.43^{+0.03}_{-0.03}$
J2250.1+3825	$1.91^{+0.84}_{-0.66} \times 10^{-9}$	8.83×10^{-10}	0.46	2.72	$1.90^{+0.06}_{-0.05}$
J0627.0-3529	$2.85^{+1.12}_{-0.80} \times 10^{-9}$	1.01×10^{-9}	0.37	2.77	$1.87^{+0.04}_{-0.05}$
J0622.4-2606	$5.41^{+3.71}_{-2.69} \times 10^{-9}$	8.75×10^{-10}	7.75	2.12	$1.63^{+0.14}_{-0.11}$
J2131.5-0915	$4.21^{+3.26}_{-2.08} \times 10^{-9}$	6.00×10^{-10}	4.26	2.27	$1.64^{+0.12}_{-0.12}$
J0115.8+2519	$2.08^{+1.03}_{-0.78} \times 10^{-9}$	9.38×10^{-10}	0.80	2.54	$1.89^{+0.07}_{-0.06}$
J0757.0+0956	$2.65^{+6.98}_{-2.46} \times 10^{-9}$	8.99×10^{-10}	59.98	2.25	$1.62^{+0.93}_{-0.61}$
J0958.6+6534	$4.25^{+6.44}_{-3.71} \times 10^{-9}$	5.73×10^{-10}	46.64	2.35	$1.35^{+0.68}_{-0.34}$
J2329.2+3754	$1.99^{+1.25}_{-0.95} \times 10^{-9}$	6.34×10^{-10}	2.19	2.19	$1.81^{+0.11}_{-0.08}$
J1436.8+5639	$1.39^{+0.89}_{-0.68} \times 10^{-9}$	3.18×10^{-10}	1.05	2.54	$1.79^{+0.10}_{-0.09}$
J1442.8+1200	$1.54^{+0.84}_{-0.67} \times 10^{-9}$	3.57×10^{-10}	0.49	2.69	$1.81^{+0.08}_{-0.07}$

^a $E^2 dN/dE$ at (1+z) TeV

^b Uncertainty in Γ at (1+z) TeV




Figure 26. Source Images and SEDs for objects in Table 1 – part 1. Left: Compilations of the 3FGL (red/orange) and 2FHL (blue/light blue) fluxes. Also shown are the broken power-law fits (dashed lines) and associated uncertainty (grey region). Right: 1–100 GeV count maps for the sources on the left.




Figure 27. Source Images and SEDs for objects in Table 1 – part 2. Left: Compilations of the 3FGL (red/orange) and 2FHL (blue/light blue) fluxes. Also shown are the broken power-law fits (dashed lines) and associated uncertainty (grey region). Right: 1–100 GeV count maps for the sources on the left.

- Böttcher, M., Reimer, A., Sweeney, K., & Prakash, A. 2013, *ApJ*, 768, 54
- Broderick, A. E., Chang, P., & Pfrommer, C. 2012, *ApJ*, 752, 22
- Broderick, A. E., Pfrommer, C., Puchwein, E., & Chang, P. 2014a, *ApJ*, 790, 137
- Broderick, A. E., Pfrommer, C., Puchwein, E., Chang, P., & Smith, K. M. 2014b, *ApJ*, 796, 12
- Broderick, A. E., Tiede, P., Shalaby, M., Pfrommer, C., Puchwein, E., Chang, P., & Lamberts, A. 2016, in prep.
- Cerruti, M., Zech, A., Boisson, C., & Inoue, S. 2015, *MNRAS*, 448, 910
- Chang, P., Broderick, A. E., & Pfrommer, C. 2012, *ApJ*, 752, 23
- Chang, P., Broderick, A. E., Pfrommer, C., Puchwein, E., Lamberts, A., & Shalaby, M. 2014, *ApJ*, 797, 110
- Chen, W., Buckley, J. H., & Ferrer, F. 2015, *PhRvL*, 115, 211103
- Dermer, C. D., Cavadini, M., Razzaque, S., Finke, J. D., Chiang, J., & Lott, B. 2011, *ApJL*, 733, L21
- Dolag, K., Kachelriess, M., Ostapchenko, S., & Tomàs, R. 2011, *ApJL*, 727, L4
- Domínguez, A., Finke, J. D., Prada, F., Primack, J. R., Kitaura, F. S., Siana, B., & Paneque, D. 2013, *ApJ*, 770, 77
- Duplessis, F., & Vachaspati, T. 2017, *ArXiv e-prints*
- Durrer, R., & Neronov, A. 2013, *A&ARv*, 21, 62
- Fossati, G., Maraschi, L., Celotti, A., Comastri, A., & Ghisellini, G. 1998, *MNRAS*, 299, 433
- Ghisellini, G., Celotti, A., Fossati, G., Maraschi, L., & Comastri, A. 1998, *MNRAS*, 301, 451
- Gould, R. J., & Schröder, G. P. 1967, *PhRv*, 155, 1408
- H. E. S. S. Collaboration et al. 2014, *A&A*, 562, A145
- Hopkins, P. F., Richards, G. T., & Hernquist, L. 2007, *ApJ*, 654, 731
- Kandus, A., Kunze, K. E., & Tsagas, C. G. 2011, *PhR*, 505, 1
- Kneiske, T. M., Bretz, T., Mannheim, K., & Hartmann, D. H. 2004, *A&A*, 413, 807
- Lamberts, A., Chang, P., Pfrommer, C., Puchwein, E., Broderick, A. E., & Shalaby, M. 2015, *ApJ*, 811, 19
- Long, A. J., & Vachaspati, T. 2015, *JCAP*, 9, 065
- McNamara, B. R., Nulsen, P. E. J., Wise, M. W., Rafferty, D. A., Carilli, C., Sarazin, C. L., & Blanton, E. L. 2005, *Nature*, 433, 45
- Miniati, F., & Elyiv, A. 2012, *ArXiv e-prints*
- Neronov, A., Semikoz, D., Kachelriess, M., Ostapchenko, S., & Elyiv, A. 2010, *ApJL*, 719, L130
- Neronov, A., & Semikoz, D. V. 2009, *PhRvD*, 80, 123012
- Neronov, A., Semikoz, D. V., Tinyakov, P. G., & Tkachev, I. I. 2011, *A&A*, 526, 90
- Neronov, A., & Vovk, I. 2010, *Sci*, 328, 73
- Nulsen, P. E. J., McNamara, B. R., Wise, M. W., & David, L. P. 2005, *ApJ*, 628, 629
- Pfrommer, C., Chang, P., & Broderick, A. E. 2012, *ApJ*, 752, 24
- Prokhorov, D. A., & Moraghan, A. 2016, *MNRAS*, 457, 2433
- Puchwein, E., Pfrommer, C., Springel, V., Broderick, A. E., & Chang, P. 2012, *MNRAS*, 423, 149
- Pushkarev, A. B., Kovalev, Y. Y., Lister, M. L., & Savolainen, T. 2009, *A&A*, 507, L33
- R Core Team. 2014, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria
- Ramberg, J. S., & Schmeiser, B. W. 1974, *Commun. ACM*, 17, 78
- Salamon, M. H., & Stecker, F. W. 1998, *ApJ*, 493, 547
- Schlickeiser, R., Ibscher, D., & Supsar, M. 2012, *ApJ*, 758, 102
- Schlickeiser, R., Krakau, S., & Supsar, M. 2013, *ApJ*, 777, 49
- Shabala, S. S., Ash, S., Alexander, P., & Riley, J. M. 2008, *MNRAS*, 388, 625
- Sironi, L., & Giannios, D. 2014, *ApJ*, 787, 49
- Su, S. 2007, *Journal of Statistical Software*, 21
- . 2016, *GLDEX: Fitting Single and Mixture of Generalised Lambda Distributions (RS and FMKL) using Various Methods*
- Takahashi, K., Mori, M., Ichiki, K., & Inoue, S. 2012, *ApJL*, 744, L7
- Tavecchio, F., Ghisellini, G., Bonnoli, G., & Foschini, L. 2011, *MNRAS*, 414, 3566
- Tavecchio, F., Ghisellini, G., Foschini, L., Bonnoli, G., Ghirlanda, G., & Coppi, P. 2010, *MNRAS*, 406, L70
- Taylor, A. M., Vovk, I., & Neronov, A. 2011, *A&A*, 529, A144
- Tiede, P., Broderick, A. E., Shalaby, M., Pfrommer, C., Puchwein, E., Chang, P., & Lamberts, A. 2016, in prep.
- Urry, C. M., & Padovani, P. 1995, *PASP*, 107, 803
- Wakely, S. P., & Horan, D. 2008, *International Cosmic Ray Conference*, 3, 1341
- Zacharias, M., & Wagner, S. J. 2016, *A&A*, 588, A110